

The Importance of Artificial Intelligence and Data for the Telecommunications Industry and the FCC

**Report of the FCC's
Technological Advisory Council
Working Group on Artificial Intelligence and
Computing
(FCC TAC AIWG)
January 14th, 2021**




Illustration Source: Techmesquare – Top 10 Interesting facts about Artificial Intelligence

Table of Contents

(i)	The FCC TAC Artificial Intelligence Working Group Team Members
1	Executive Summary
2	Introduction and Background
3	AI/ML Focus Areas of Relevance to FCC
3.1	AIWG Objectives
3.2	Considerations for Achieving the AIWG's Objectives
4	Status of AI/ML Technology: What it Takes to Deliver on AI
4.1	Impacts of AI/ML on an Organization's Digital Transformation
4.2	The Importance of Data
4.3	An Example - Addressing an FCC Priority: Closing the Digital Divide
4.4	Practicing the Safe Use of AI
4.5	Leveraging Federally Funded Initiatives
5	Key Findings, Conclusions, and Recommendations
5.1	Inclusion of Considerations for AI/ML in the FCC's Strategic Plan
5.2	Creating an AI/ML Data Task Force Within the FCC
5.3	Implementing a Procurement Process for Inclusion of AI/ML Capabilities in the FCC's National Broadband Map (NBM)
5.4	Considerations for the Safe Use of AI/ML by the FCC and the Telecommunications Industry
5.5	Initiating Pilot Projects
6	General References
7	Appendix A. AI/ML Landscape & Annotated Bibliography
8	Appendix B: AI/ML in Federally Funded Initiative
8.1	Long Term and Targeted Investments in Research and Development

- 8.2 Human-AI Interactions and Collaboration.
- 8.3 Ethical, Legal, Societal Implications of AI
- 8.4 Safety & Security of AI Systems
- 8.5 Data Sets for Training & Testing AI Systems
- 8.6 Measuring & Evaluating AI System Performance
- 9 Appendix C: DARPA’s Spectrum Collaboration Challenge: A Review
 - 9.1 Introduction
 - 9.2 An Experimental Testbed – The Colosseum
 - 9.3 The Spectrum Collaboration “Game”
 - 9.4 Impediments to Collaboration
 - 9.5 Team Performance and Overall Results
 - 9.6 Extending DARPA’s Work for Spectrum Sharing for Commercial Users
 - 9.7 Moving Beyond Bi-lateral Bargaining
- 10 Appendix D: Data Management
- 11 Appendix E: Safe Use of AI
- 12 Appendix F: Uses of AI/ML in the Creation of a Broadband Map
- 13 Appendix G: Pilot Project Descriptions
 - 13.1 Pilot Project on Testing Lab
 - 13.2 Pilot Project to Improve Innovation Using AI/ML Benchmarking
 - 13.3 Pilot Project on AI to Aid Enforcement
 - 13.4 Pilot Project on Self Organizing Networks
 - 13.5 Pilot Project - Employing AI to Assess RF Interference and Promote 5G Open Radio Access Networks
- 14 Appendix H: List of Presentations to the FCC TAC AIWG

(i) **The FCC Technological Advisory Council Working Group on Artificial Intelligence and Computing Team Members and FCC Liaisons**

2020 FCC TAC AIWG Team Members

Chairs: Lisa Guess, Cradlepoint
Adam Drobot, OpenTechWorks, Inc.

FCC Liaisons: Michael Ha, Mark Bykowsky, Monisha Ghosh,
Martin Doczkat, Robert Pavlak, Chrysanthos Chrysanthou,
Gulmira Mustapaeva

- | | |
|-----------------------------------|---|
| • Shahid Ahmed, Independent | • Nageen Himayat, Intel |
| • Sujata Banerjee, VMware | • Steve Lanning, Viasat |
| • Nomi Bergman, Advance | • Gabriel Lennon, Intern Univ of Colorado |
| • William Check, NCTA | • Kevin Leddy, Charter |
| • Brian Daly, ATT | • Anne Lee, Nokia |
| • Adam Drobot, OpenTechWorks | • Brian Markwalter, CTA |
| • Jeffrey Foerster, Intel | • Lynn Merrill, NTCA |
| • James Goel, Qualcomm | • Jack Nasielski, Qualcomm |
| • Lisa Guess, Cradlepoint | • Michael Nawrocki, ATIS |
| • Russ Gyurek, Cisco | • Dennis Roberson, entigenlogic |
| • Dale Hatfield, Univ of Colorado | • Marvin Sirbu, SGE |
| • Stephen Hayes, Ericsson | • David Tennenhouse, VMware |
| • Mark Hess, Comcast | |

1 Executive Summary

The Federal Communications Commission Technological Advisory Council Working Group on Artificial Intelligence and Computing (AIWG) addressed the impacts that wide scale deployment and adoption of Artificial Intelligence methods and techniques may have on the Nation's Networks. The AIWG emphasis was on capturing the promise and benefits of AI as an integral part of the Nation's Telecommunications ecosystem. At the same time the AIWG also examined approaches to promote practices that lead to the safe uses of AI and minimize common pitfalls. The WG recognized the large investment in AI technologies by Industry and the Federal Government.¹ It also took into account the considerable momentum of the Basic and Applied Research Community in maturing AI technologies, and the significant penetration in the marketplace of commercial services and applications that rely on AI, including those offered (and used internally) by the Telecommunications sector. The considerations included impacts on:

- The FCC
- Network Operators and Service Providers
- Consumers
- Industry, Commercial Enterprises, and Government
- The Nation's Networks

To conduct its deliberations the AIWG formed two Sub-Working Group teams (SWGs) for addressing this year's charter. The first of these, led by Nageen Himayat from Intel, focused on Federal Investments in AI and met on a weekly basis for discussions, presentations from subject matter experts, and in-depth technical explorations where warranted. The second SWG, led by Nomi Bergman from Advance and David Tennenhouse of VMware, focused on Safe uses of AI and met as part of the main AIWG. Each of the SWGs teams also held a limited number of team meetings to draft findings and recommendations. The main AIWG met on a weekly basis for sessions that were devoted to discussions and deliberations, and presentations from a wide variety of experts on AI research, specific aspects of technologies necessary for AI, and applications of AI. In total the AIWG heard 26 presentations involving 38 individuals, several of whom returned for in depth technical discussions.

The term Artificial Intelligence is identified with many different methods and techniques and specialized application areas (such as robotics, machine vision, and natural language processing among others.) The term also evokes the not un-common perception of computers being able to exhibit traits to think like human beings and perhaps surpass human

¹ On November 17, 2020, the Office of Management and Budget (OMB) published a memorandum that provides guidance to all Federal agencies regarding technologies and industrial sectors that employ artificial intelligence (AI). The memorandum establishes policy considerations and goals that should guide, to the extent possible by law, regulatory and non-regulatory approaches to AI applications developed and adopted by the private sector. See, <https://www.whitehouse.gov/wp-content/uploads/2020/11/M-21-06.pdf>.

abilities. Such a capability is referred to as General AI. At this point in time such a capability does not exist, and it is uncertain when it will emerge and reach maturity. In general, the AIWG took a broad view of AI and considered the utility of AI methods where they could create value by performing tasks that are hard for human beings to do. This includes the automation of routine processes, the ability to handle large numbers of simple transactions at scale, optimization of systems for performance and efficiency, and to conquer the complexity of problems that have been too hard to untangle in the past. Most important is AI as a source of innovation, that will allow us to create better business models, more efficient and flexible manufacturing methods, new products and services, and greatly improve the processes in our enterprises and in our personal lives.

The AIWG focused much of its attention on Machine Learning (ML). As an aspect of Artificial Intelligence, Machine Learning holds much promise, and it is hoped that the methods associated with ML will make profound contributions to solve some of the intransigent problems in networking. These include the efficient use of spectrum, the automation to carry the rapidly growing number of network connections projected by analysts, and innovative use of evolving network capabilities in important application areas. At the same time the AIWG also identified issues surrounding the broad use of AI deserving further attention.

The AIWG found that AI methods and techniques have broad applicability across the Telecommunications Industry and are important to how the FCC may carry out its mandates in the future. Two key aspects of ML are the importance of Data and access to Computing and Storage resources. The first of these has implications for the role that the FCC might play in the AI ecosystem and the second has fundamental implications for Future Network Architectures. Lastly in its deliberation the AIWG took into account the maturity of AI technologies as applied to Telecommunications, projection of timescale for impact, the need for resources, and the alignment with the FCC's primary strategic objectives:

- Closing the Digital Divide
- Promoting Innovation
- Protecting Consumers & Public Safety
- Reforming the FCC's Processes

In its conclusions the AIWG arrived at five recommendations, the first two of which can be thought as overarching with long term impact on the FCC and the Telecommunications ecosystem and three that address important near-term priorities:

The FCC TAC AIWG has identified five recommendation areas:

1. **“Unlock transformational change”** - The incorporation of considerations for Artificial Intelligence in the FCC Strategic Plan.
2. **“To build knowledge, unleash the Data”** - The creation of a Task Force to address how the FCC can best address important aspects of Data governance and curation for AI/ML applications to serve its internal needs, and those of industry and the public.
3. **“Cast a wide net”** - Develop a plan and strategy for designing, developing, deploying, operating, and maintaining a Broadband Map that takes advantage of the best technologies and capabilities available.
4. **“Keep humans in control of the loop”** - Policies and approaches to ensure the safe use of Artificial Intelligence as it impacts the Nation’s Networks, communication needs, and important applications.
5. **“Get your feet wet”** - Develop the FCC’s capability for extracting value from Artificial Intelligence in solving issues and problems that come before the FCC by conducting pilot projects with near term return.

The AIWG White Paper provides further insight into the recommendations and the rationale behind them.

2 Introduction and Background

Artificial Intelligence as a term describes computational techniques consisting of multiple major branches each focused on different methods, approaches, or areas of application. Since the term “Artificial Intelligence” was coined in the mid -1950s there have been successive waves of deployment of applications that are now routine (such as Optical Character Recognition) or have been widely adopted (among others in areas such as machine vision, speech recognition, natural language processing, knowledge systems, classification systems, and applications including: noise cancellation; search; customer relationship management, chatbots for customer facing services; facial recognition; gaming; control systems; and analysis of high dimensional data).

Within the family of AI techniques, Machine Learning (ML) and its sub-branches have recently seen a dramatic surge of investment, activity, and popularity across multiple domains including those that may be of interest to the FCC, with significant impact now, in the near-term, and potentially profound impact in the long term.

The various types of ML hold significant promise for dealing with the complex problems that are important to network planning and design, control and operation of networks, the management of resources such as spectrum, for the technologies used as the building blocks of network architectures, components, devices, and applications, and lastly for customer services and automation of customer interactions through chatbots.

With the continued growth of demand for digital traffic, networks now connect over 4 Billion people on the planet and 10's of billions of devices. ML as a tool set is important because the increasing level of complexity is outpacing the ability of well-established methods in reliably delivering solutions and ML offers a potential approach for conquering that complexity. It is particularly important for machine-to-machine communications characteristic of Industrial uses, the Internet of Things, and the increased heterogeneity of digital traffic on the Network.

The sheer scale and new exacting requirements for network performance are compounding the problem. There is a widely recognized need for a higher level of automation of network functions to deliver efficient and in some sense near optimal solutions. Examples can be found in areas such as spectrum sharing; network access control; security and privacy, management of 5G Networks; operational aspects of wireless network densification, and in specific new capabilities such as network slicing. AI is not just a tool; it has the potential of being a game changer.

It is important to say that AI methods are far from achieving what is referred to as “General AI”– that is the ability to reason like human beings and to be able to make the kind of generalizations, that human beings are capable of, or common-sense reasoning in the face of ambiguity or previously un-encountered conditions. There are however techniques and methods within AI, such as Reasoning, Rule Based, and Knowledge Based Systems (Narrow AI) that significantly exceed human performance. These are typically single threaded in nature and focused on solving a single issue or action at a time.

The general approach of this AI Working Group was the following:

- Develop use cases and important issues where AI has a potential impact on the FCC, the national telecommunications ecosystem, and can be an enabler for significant new economic activity.
- Study issues and use cases leveraging industry experts, analysts, government organizations, and the research community and gather relevant background data
- Categorize and prioritize use cases and issues (based on relevance to the FCC, the industry, or end users, likely level of impact, and time frame to maturity)
- Understand use cases in context of AIWG Charter and the objectives

3 AI/ML Focus Areas of Relevance to FCC

Conducting this year’s AIWG benefited greatly from the prior four years when the WG had studied “Future Network Technologies”, “Computational Power Stress on the Network” motivated by the increasing change in use patterns brought on by increasing machine - machine traffic and cloud computing, and the increased use of software defined functionality in Network control and management. In approaching this year’s objectives, the WG considered the FCC Strategic Plan and its main objectives, the trend in Network patterns that saw the big shift to Industrial uses and the projection of domination by machine-machine traffic. Lastly the WG considered the adoption of AI technologies by the Telecommunications sector - which is considered as one of the lead sectors in AI, and by the proliferation of AI based products in the marketplace.

As part of its deliberations the WG also addressed the criteria for evaluating the aspects of AI to concentrate on. These included the maturity of AI technologies, the likely impact on important applications, the importance to the FCC’s processes, and the likely timescale on which the AI technologies may be important. In the last case we prioritized items that are likely to provide benefits on a short timescale but also included strategic considerations that are important to build capacity at the FCC to take advantage of advances in AI Technologies. The AIWG also considered important trends such as the growth of machine-machine connections on the Network, shown in Figure 3.0.1, and the evolving hierarchy of Networking and Computational Architectures, shown in Figure 3.0.2.




Figure 3.0.1: The shifting pattern of Network connections - AI is considered as an important factor in the growing demand for computational and storage resources in Industrial Applications

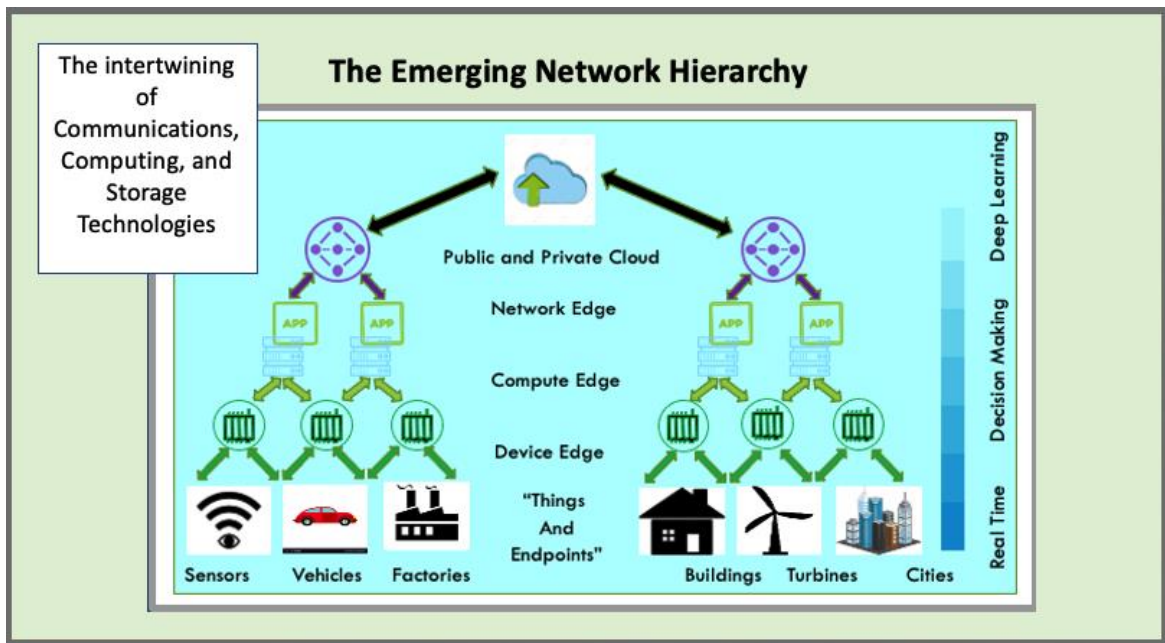


Figure 3.0.2: The use of AI/ML technologies is having a considerable impact on Network Architecture and illustrates the deep intertwining between communications, computing, and storage.

3.1 AIWG Objectives

The Artificial Intelligence (AI) and Computing Working Group was chartered to analyze the ability of AI to improve the performance of telecommunications networks and the services enabled by these networks. The AI Working Group prioritized the following objectives that had high relevance for the FCC.

- **Objective 1:** Identify how to Leverage Federally Funded Initiatives. Identify how can the results from recent programs in AI for spectrum and networking, such as the DARPA Spectrum Collaboration Challenge (SC2) and the NSF/Intel joint solicitation on Machine Learning for Wireless Networking Systems (MLWiNS), be leveraged for real-world systems and applications and for investigating new applications?
- **Objective 2:** Detail the datasets required to drive AI innovation. AI relies on curated and labeled data sets being available for algorithm development and testing: what should the parameters of such data sets be?
- **Objective 3:** Use AI to Extract Meaningful Information from Existing Data. How can AI be used to extract meaningful information from data that are either already available (e.g., from the Measuring Broadband America (MBA) program) or may become available.
- **Objective 4:** Identify the Risks and Mitigation Strategies in the Use of AI. As legitimate applications of AI start proliferating, what risks should be evaluated and what AI tools exist or should be developed to identify and mitigate harms that might arise from the proliferation of AI?

3.2 Considerations for Achieving the AIWG's Objectives

As noted, the FCC's Strategic Plan has, as one of its primary objectives, the adoption of policies and procedures that promote a competitive, vibrant, and innovative telecommunications marketplace. A competitive marketplace is important because, under such conditions, firms are induced to make decisions that help drive resources to their most efficient and highest valued use. A vibrant market is one in which participants are able and willing to dynamically respond to changes in those market conditions. Finally, an innovative market is one in which firms adopt new, and more cost-effective methods of providing existing services and/or provide entirely new services. Taken together, markets that are characterized by these three features are the major source of sustainable growth for the economy.

The FCC directed the AIWG to explore four specific areas of interest. Each of these areas touches on the FCC's ability to satisfy its responsibility to Congress and the American people for ensuring a vibrant competitive marketplace. The first area touches on the FCC's goal of promoting the introduction of new technologies and services. It involved examining the progress that federally funded research has made in exploring the ability of AI to improve the efficiency with which spectrum is used. One such research program (i.e., DARPA) examined the spectrum sharing performance properties of a highly decentralized

system in which an administrator (e.g., FCC) plays a very limited spectrum management role. In this program, for example, the administrator's customary role of identifying frequency and bandwidth allocations is replaced entirely by the decentralized spectrum access decisions made by users in response to performance and other rules established by the administrator.

Another area of examination is directly related to the FCC's goal of "closing the digital divide." In particular, the AIWG was directed to explore whether AI, when applied to existing broadband coverage data, can be used to provide a more accurate estimate of the download and upload speeds broadband providers offer their customers. The AIWG went further to examine whether technological and other advances allow for a completely new way of gathering broadband data and, with the assistance from AI, aid in the development of more accurate broadband maps. The goal is not just the maps, but an interactive system that can aid stakeholders in making better decisions.

Yet another area of AIWG's examination falls squarely within another strategic goal of the FCC, namely the goal of protecting consumers and promoting public safety. It is well recognized that while AI has the potential to significantly improve the welfare of individuals, when used in the wrong way it also has the ability to impose substantial harm upon society. The AIWG explored the risks associated with the proliferation of AI, and identified a set of rules and conditions that, when satisfied, would ameliorate the harm society may experience from the misuse of AI.

Finally, the last area of examination is related to the FCC's goal of reforming its operations and processes. Here, the AIWG explored the importance of data in leveraging advances in AI. In particular, the AIWG examined the usefulness of having the FCC develop machine readable, well-curated data sets for use by itself, industry participants, and academic researchers. The AIWG examined the usefulness of initiating select pilot projects designed to promote the FCC's goal of ensuring a competitive, vibrant, and innovative telecommunications marketplace.

4 Status of AI/ML Technology: What it Takes to Deliver on AI

While we observe that AI may be thought of as a tool to enhance the Nations Networks and the processes of organizations that manage or facilitate the development and enhancement of these networks, this new tool has the potential to be a disruptive force. Therefore, careful thought is needed to drive investments that properly develop AI. These investments range from basic research to actionable items.

4.1 Impacts of AI/ML on an Organization's Digital Transformation

(Strategic AI Plan Team: Mark Bykowsky, Adam Drobot, Lisa Guess, Mark Hess, Mike Nawrocki)

Digital transformation is a multi-faceted process by which an organization employs electronic tools, systems and devices (so-called “digital technologies”) that collects, stores, and processes data with the objective of generating long lasting efficiency-enhancing and other value-generating changes in its operations.² The process of digital transformation is occurring in both the private and public sectors, though to differing degrees. In the public sector, digital transformation is driven by the need to do more with less, improve transparency, leverage data and analysis in response to the growing complexities of policy issues, and, in general, enhance the level of engagement and trust that citizens have with and in the public sector.

AI and ML are key elements in the digital transformations occurring in many public sector organizations. According to a recent survey conducted by Deloitte involving 110 public sector executives, 57% of early adopters of AI believe that AI is “very” or “critically” important in their organization’s success today, and 74% of all respondents believe it will be in the next two years.³ Early adopters within the public sector see AI not as a substitute for humans but rather as a tool to enable workers to be more creative (e.g., chatbots and virtual assistants). Such a perspective validates Alfred North Whitehead’s statement that “Civilization advances by extending the number of important operations which we can perform without thinking of them.” Among the most important use cases of AI in the public sector involves improving quality control (e.g., detecting defects) and workforce management (e.g., recruiting and training).

² Some authors view the process of “digital transformation” as comprising of three distinct steps. The first step, sometimes referred to as “digitization,” involves converting information into digital (e.g., machine-readable) form. The objective of this step is to improve the accessibility of the information. The second step, sometimes referred to as “digitalization,” involves automating the organization’s processes in an effort to lower its costs and improve productivity. The third step, sometimes referred to as “digital transformation,” involves applying digital tools to the organization’s operations in an effort to increase both the value that customers place on its services and products, and to provide completely new and innovative products and services.

³ Deloitte Insights: Government Executives on AI: Surveying How the Public Sector is Approaching an AI-enabled Future. See <https://www2.deloitte.com/global/en/insights/industry/public-sector/ai-early-adopters-public-sector.html>.

Regardless of whether the organization is public or private, achieving real and sustainable digital transformation requires developing an overarching strategy with an “organization wide” perspective that focuses on outcomes related to the organization’s mission. Among other things, it requires the ability to identify, attract, and engage talent across the organization. As it applies to AI, sustainable digital transformation requires that the organization seek additional and sometimes new skill sets that work with and are complementary to AI technologies. In addition to data scientists and software and other engineers, it involves developing expertise in project management, including individuals with expertise in identifying private sector suppliers of the necessary inputs and skills sets. Indeed, because of budget and other constraints, an organization’s ability to reach sustainable digital transformation may require working with other public sector organizations. That requirement involves having the organization possess the human capital skill in developing a strong, positive relationship with other stakeholders in other organizations. Finally, key to digital transformation success is the careful identification of use cases. Here, the organization needs skilled subject matter experts to guide it in selecting the areas in which AI can improve the organization’s performance and its ability to achieve its strategic goals. Subject matter experts have the responsibility to ensure that the use cases support and are consistent with the objectives and goals included in the FCC’s Strategic Plan. Based upon the subject matter expertise of the working group’s members, Appendix G includes a list of use cases which, upon completion, would greatly assist the FCC in reaching the goals and objectives included in its strategic plan. Importantly, the AIWG believes that the completion of each one of these use cases, referred to as “pilot projects” for purposes of the White Paper, would serve to catalyze innovation and stimulate U.S. economic growth and prosperity.

4.2 The Importance of Data

(The Data Team: Sujata Banerjee, Mark Bykowsky, James Goel, Stephen Hayes, Nageen Himayat, Anne Lee)

Data has become an increasingly important resource in the economy. Its importance stems from the value of the information it provides its users. The various ways in which data are used in the economy is vast. For example, data are often used to provide evidence that assists in resolving disputes. It is often used by businesses to measure the effectiveness of a given strategy. Scientists employ data to assist in testing theories about the fundamental forces of nature. Medical researchers use data to better understand the origin, transmissibility, the mutation capability, and the proper treatment of deadly viruses.

Data plays a critical role in the performance and operations of the communications service sector. For example, the demand that users place on a telecommunications network at a given moment in time often requires the service provider to, in an effort to minimize any negative quality of service effects, reallocate some of its network resources. One approach to minimizing quality of service effects involves attempting to predict changes in user demand before they take place. Those predictions can then be used to initiate the necessary

changes in the network's resources. Without data, the service provider would find it extremely difficult to improve the customer experience and to increase profitability.

The growth in the importance of data has corresponded to advances in analytics, including ML. For example, ML now affords wireless service providers the opportunity to make more accurate predictions regarding changes in user demands. A necessary requirement for the development of more accurate predictions in this and other situations is the existence of, and access to, well-curated and cleansed data.[1] Without such data, the resulting model(s) will not be able to uncover the hither-to-fore unknown patterns that yield better predictions which could lead to improved quality of service in telecommunications networks. The identification of such patterns constitutes *knowledge*, the leveraging of which can greatly improve the performance of not only telecommunication networks but can, when broadly obtained and properly leveraged throughout the economy, lead to a substantial improvement in the human condition.

An important goal of the FCC is to ensure that its decisions are fact-based, and data driven. Indeed, it has stated that data underpins every one of its activities.[2] Further, it has the responsibility of managing numerous important databases (e.g., broadband service data).[3] The FCC's ability to effectively collect and manage data is critical not only because of its reliance on data in its decision making and other processes, but also because its decisions have important economic effects far beyond the telecommunications and electronics equipment sectors. The AIWG believes that recent advances in ML and other analytical techniques creates an opportunity for the FCC to harvest additional and important *knowledge* from the databases it currently manages and may manage in the future. Further, the *knowledge* obtained from AI-based analytics are likely to significantly impact all of the FCC's high level strategic goals which include promoting an innovative telecommunications market, expanding service to the under-served, protecting consumers, and advancing internal FCC processes as highlighted earlier. Taking advantage of this opportunity requires careful analysis of the current state of data affairs and management policies at the FCC.

It also requires building an AI/ML Data System, defined as a set of resources and processes that are needed to train a ML-based algorithm or model. Figure 4.2.1 below (derived from [7]) describes the various components of such a system. The size of each box roughly corresponds to the code/infrastructure that is needed for the identified task. There are several notable features of an AI/ML Data system. First, the process of collecting and verifying data, which includes such things as data collection, curation, labeling and storage, are not only two essential tasks of an ML system, but they also require a considerable amount of resources for their completion.[4] Second, only a small fraction of the overall code is actually attributed to the *ML Algorithms*.

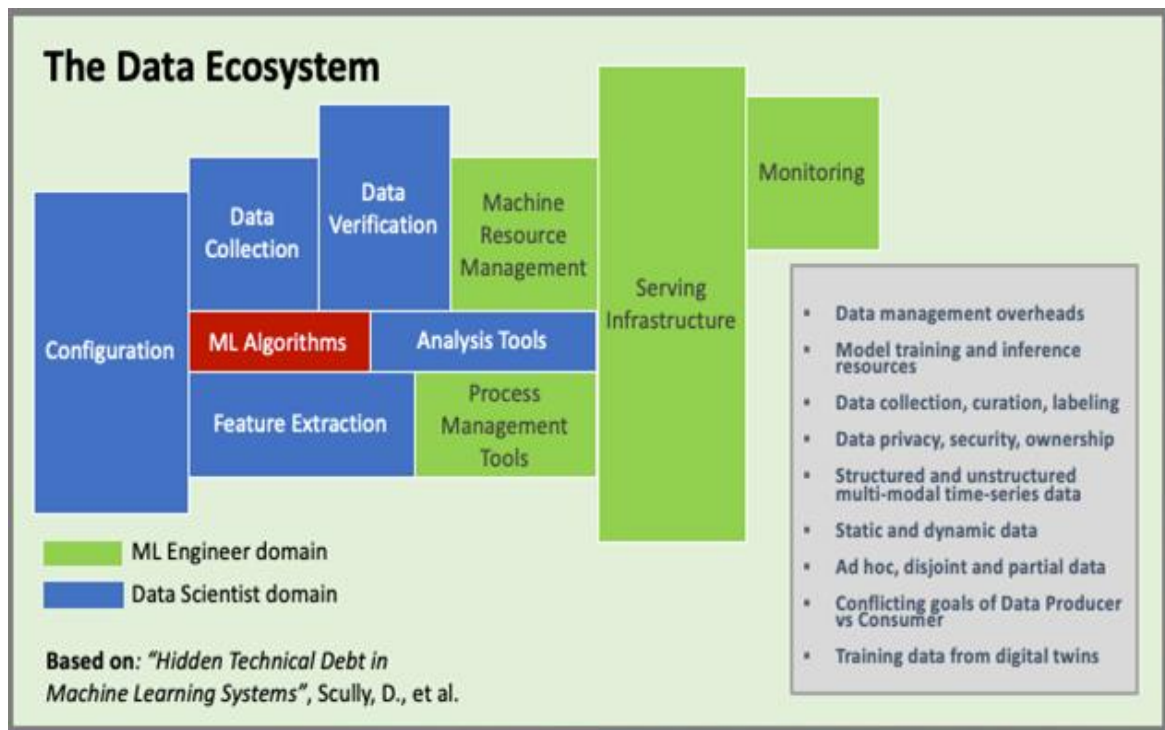


Figure 4.2.1: In a Data Ecosystem ML Algorithms are Surrounded by a Complex Set of Technical Domains

The AIWG firmly believes that the benefits that flow from the FCC’s possible decision to develop and maintain well-curated data sets, and the construction of an AI/ML Data System, far exceed the cost incurred from doing so.[5] In addition to improving policy decisions, well-curated, machine-readable data will be used by industry and academic researchers to build a variety of AI/ML models the result of which would promote competition and spur innovation in the telecommunications sector. For example, the existence of well-curated data introduces an intriguing new possibility. Assume for the moment that a standardized data set involving the performance properties of a set of telecommunications network components existed. AI/ML techniques can be employed to develop, based upon that data, a “benchmark” (i.e., point of reference) regarding the performance of each of those components over a wide variety of operational environments. A comparison of that benchmark with the actual performance of a firm’s chosen network component would allow the operator to make more informed purchase decisions regarding which inputs to include in its network. The development of such a benchmark would spur competition and innovation in the network component market and improve the customer experience.

In many instances, the absence of well-curated data is not the result of technical or financial matters, but rather because of the private welfare maximizing decisions of their owner. In many cases such owners can have “too great” of an incentive from society’s point of view not to share their data in that the gains that society obtains from sharing the data exceed the

cost the data owner incurs from that sharing. The resulting effect is a diminished level of innovation in the U.S. economy broadly, and a reduction in consumer and producer welfare in the telecommunications sector. Importantly, the FCC can play an important role in correcting the misalignment of interests between society and the owners of the data. Among other things, the FCC could encourage data owners to adopt new and innovative privacy preserving techniques to their data.[6] Second, the FCC could, by organizing and participating in meetings involving interested parties, facilitate solutions to the various hurdles that exist involving the sharing of data held by the owner of the data. Third, the FCC could begin to explore the feasibility of promoting the development of a data exchange. Appendix D presents additional material related to data management and the use of AI/ML tools.

End Note [1] While data come in a wide variety, they all share some underlying, somewhat abstract characteristics. For example, data contains: (1) information that has meaning to the user; (2) attributes that serve to differentiate one element of data (datum) from another; and (3) information about the source of the data.

End Note [2] See <https://www.fcc.gov/reports-research/data>.

End Note [3] See <https://www.fcc.gov/general/broadband-deployment-data-fcc-form-477>. See also <https://www.fcc.gov/licensing-databases/search-fcc-databases>.

End Note [4] “Curation” is the organization, integration and presentation of data from various sources in a manner that preserves its usefulness and reuse over time.

End Note [5] The AIWG believes that the benefits from converting many of the FCC’s data sets into machine readable, well curated data sets far exceeds the costs associated with completing that transformation. In the event that the FCC believes there are better uses of its funds, the AI WG implores the FCC to explore the opportunity to share the cost of that transformation with other Federal Agencies and Departments.

End Note [6] See presentation by Rafail Ostrovsky (UCLA/Stealth Software) “Stewardship of Private Data with Cryptography.”

End Note [7] Scully, D., et. al., “Hidden Technical Debt in Machine Learning Systems,” NeurIPS 2015.

4.3 An Example - Addressing an FCC Priority: Closing the Digital Divide

(Broadband Mapping Team: Lynn Merrill, Steve Lanning, Bill Check, and Marvin Sirbu)

The 116th Congress, in 2020 passed and subsequently funded the Broadband Deployment Accuracy and Technological Availability Act, at the level of \$65 million. The Broadband DATA Act provides the FCC with the mechanisms to improve the FCC mapping process by enhancing the collection, verification and reporting of broadband data.




Figure 4.3.1: Data on the Number of Households That Have Access to a Given Broadband Speed and the Number of Broadband Providers That Provide That Speed

Source: FCC Broadband Reports June 2014 through December 2019 as organized by TAC Member Steve Lanning

Mapping broadband availability is an important priority for the FCC as it is used to assist in policy making. There is the obvious need to know how widely broadband is deployed, how it is evolving and, what areas and locations are in need of new policies and/or support to bring broadband to every location in the US. The following can be derived from the history of 477 data.

An early version of this view was presented by a former Chairman Wheeler. His remarks remain the same that competition, even as captured by 477 data by census block, implies very little choice at the household level and therefore low levels of actual competition. This is illustrated in Figure 4.3.1. While some census blocks have three or more providers and a reasonable level of competition, most census blocks have two providers or less. As the

definition of broadband has evolved from 10 Mbps to 25 Mbps, the proportion of households with access to three or more providers at 10 Mbps in 2014 was 23.6%. In December 2019, the equivalent percent of households by census block at 25 Mbps service appears to be 30.7% or a 7% increase in the proportion of households with access to the basic defined level of broadband service (offered by three or more providers). This estimated change relies on the assumption that all households in a census block reported as covered actually have access to the maximum advertised service in the census block. This assumption is not generally true, particularly in rural areas. The degree of this difference has been documented by the state of Georgia. Figure 4.3.2 below shows the large broadband measurement effects of assuming that all households in a census block reported as covered have access to the maximum advertised service in the census.

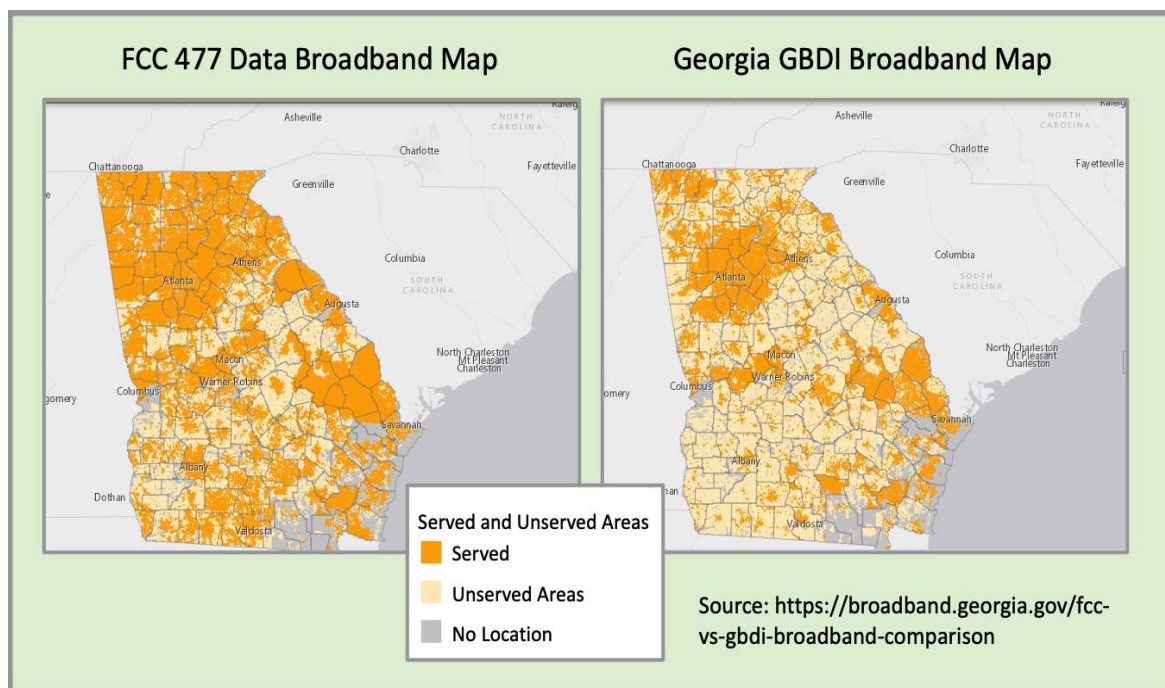


Figure 4.3.2: A Comparison of Broadband Data Measurements in Georgia

Source: <https://broadband.georgia.gov/fcc-vs-gbdi-broadband-comparison>

This has given rise to a mandate to measure the broadband fabric of the United States. For terrestrial services this means availability of a minimum level of broadband service by location. For wireless service this means detailed shape files of service availability. The FCC already has location-based support defined for RDOF Phase II. The COVID pandemic has underlined the urgency of universal broadband service and the timing of Phase II depends on solid data driven solutions that are ideally driven by AI/ML methods not only for the immediate map but also for accurate and timely data. It may be noted that the most recent 477 report is dated December 2019, but the data first became available December 2020 or about one year after the date. The data is released twice a year so the data publicly

be characterized by the *time constants*, (i.e., response times) at which they operate. These can range from milliseconds (e.g., for network control) to hours/days and all the way to days/years (e.g., for network planning). The time constants of interest can have significant implications on how AI/ML techniques are chosen and made to work safely and correctly at the required timescales. For example, some decisions in the network control sub-system will need to be made at the sub-second level. Since these models will operate faster than humans can intervene, they cannot be subject to *human-in-the-loop* control. Nonetheless, we believe that these systems can and should be designed so that humans remain *in control of the loop*.⁴

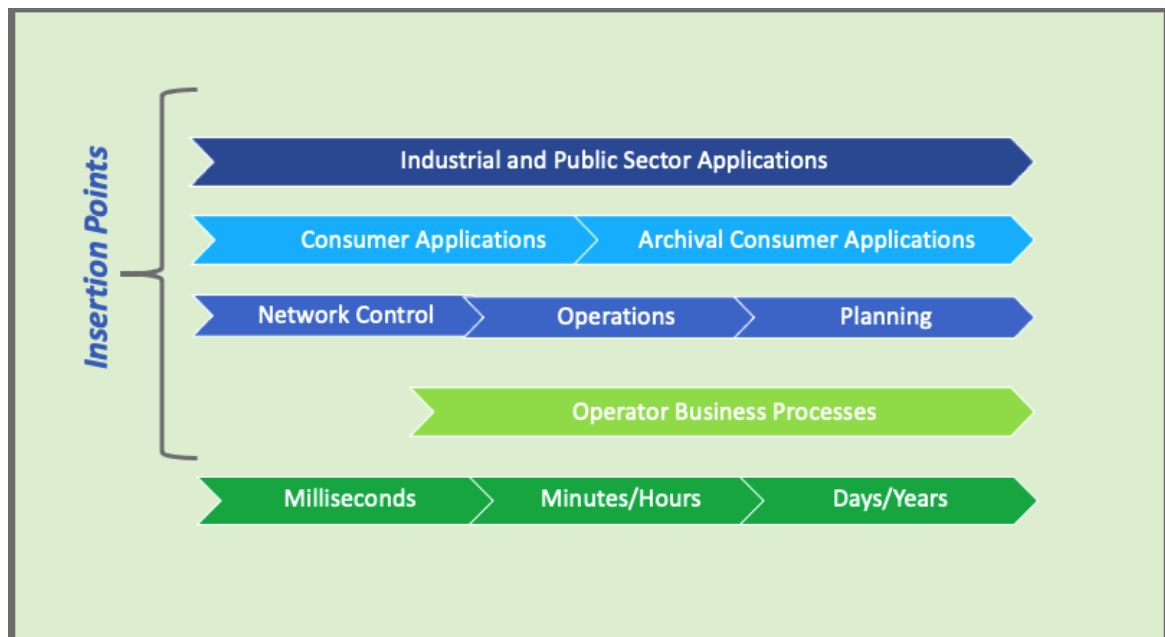


Figure 4.4.1: The insertion points for AI/ML technologies as characterized by timescale and application

For the insertion points shown in Figure 4.4.1 above, the following characteristics (or a subset) of safe AI practices are important:

- Fairness
- Transparency
- (Non) vulnerability
- Accountability and Explainability
- Robustness

⁴ At longer time constants, the throughput of ML decision-making may also preclude *human-in-the-loop* control due to the scale of the systems and the corresponding volume of ML-driven decisions.

Explanations of these characteristics along with additional details are available in Appendix E (see table in Section 11). In a broader sense Safe AI/ML also implies that its uses are Ethical, Trustworthy, and Safe and provide adequate protection for Privacy and Security.

4.5 Leveraging Federally Funded Initiatives

(Pilot Projects Team: Mark Bykowsky, James Goel, Dale Hatfield, Anne Lee, Robert Pavlak)

In an effort to enhance the growth of the U.S. economy and to maintain global leadership in Artificial Intelligence (AI), the United States Government has funded a wide variety of initiatives designed to advance research and development in Artificial Intelligence (AI). Appendix B provides a brief summary of some of the most important US Government funded initiatives that have a bearing on the provision and demand for telecommunications and information services. These initiatives fall within the following research and development areas and have relevance for the FCC.

- Long Term and Targeted Investments in Research and Development
- Human-AI Interactions and Collaboration
- Ethical, Legal, Societal Implications of AI
- Safety & Security of AI Systems
- Data Sets for Training & Testing AI Systems
- Measuring & Evaluating AI System Performance

The AIWG heard from a broad range of experts across the multiple federal agencies responsible for basic research, standards, and other initiatives with presentations from DARPA, NSF, IARPA, DHS, NIST.

Regarding the basic and applied research our general observation is that the programs we reviewed are attempting to leverage AI to its fullest and attacking important issues. The extensive research bibliography included in Appendix A, is illustrative of this coverage. While the programs have good research coverage across the diverse application of AI for telecommunication networks they are in early stages of maturity and more work is required to develop a fundamental understanding of AI solutions, the potential gains of applying AI techniques to the design of wireless networks as well as fully addressing the safety risks associated with AI-based system design.

We especially note the relevance of DARPA's spectrum challenge program to the FCC. Appendix C describes the various aspects of the program, the success and limitations of AI techniques for this challenge as well as illustrates how key components of the program are extended for FCC and commercial benefit.

5 Key Findings, Conclusions, and Recommendations

5.1 Inclusion of Considerations for AI/ML in the FCC's Strategic Plan

Driven by competitive forces and advances in data and computer science, telecommunications service providers are increasingly adopting AI and ML techniques for the purpose of furthering their digital transformation. The resulting investments are designed to reduce the cost the providers incur from providing an existing service, improve the service experience of their customers, and make possible new and innovative services. The societal and private benefits of AI and ML are not limited to the private sector. In an effort to enhance the growth of the U.S. economy and to maintain global leadership in AI, the United States Government has funded a wide variety of initiatives designed to advance research and development in AI. Such research has important uses in both the private and public sectors.

As discussed in Section 5.5, the AIWG has identified several pilot projects that involve AI and ML to explore, examine, and resolve important policy issues before the FCC. Importantly, all of these pilot projects directly support the FCC's efforts in achieving the important goals and objectives identified in its Strategic Plan. For example, one proposed pilot project involves initiating a *Request for Information (RFI)* involving the use of AI and ML in the development of more accurate broadband maps. The development of more accurate broadband maps will assist the FCC in addressing the pernicious effects of unequal broadband access across the United States. Another proposed pilot project involves exploring the use of AI and ML in developing data that accurately captures the level of interference in the RF environment within a given geographic area. Such information would enable the FCC to establish and adjust service rules in a manner that leads to a more spectrally efficient coexistence among and between spectrum users. Another proposed pilot project involves establishing a working group that explores the safe and not so safe use of AI. Finally, the FCC has the responsibility of managing numerous important databases (e.g., broadband service data). The AIWG believes that recent advances in ML and other analytical techniques creates an opportunity for the FCC to harvest additional and important *knowledge* from the databases it currently manages and may manage in the future.

Given its oversight authority over telecommunications service providers, the growing extent to which AI and ML is being employed by those and other parties over which it has authority, and the extent to which AI and ML may enable the FCC to achieve its high level strategic and policy goals, it stands to reason that AI and ML should take a prominent role in future FCC strategic plans. Considerations for AI/ML in the strategic plan should be framed to address the important priorities that the FCC sees in its mission and the possible value that AI/ML can contribute:

- *Closing the Digital Divide (Example)*
 - Improve the efficiency with which Federal Funds are used to close the digital divide by exploring the use of AI and ML for supporting investment decisions

- *Promoting Innovation (Examples)*
 - Deployment of new infrastructure to support Smart Cities
 - Adopt policies that facilitate the deployment and wide use of autonomous vehicles (possibly refocus to IOT or facilitate the development of smart cities). The examples require the continuous sharing of data, ubiquitous bandwidth, and a high degree of reliability
- *Protecting Consumer and Public Safety (Examples)*
 - Policies that facilitate the development of infrastructure (and enforcement of responsible and safe use of AI)
 - Improved web-based services for access to communications and use of information provided by the FCC
- *Reforming the FCC's Processes (Examples)*
 - Explore the use of AI and ML to expedite equipment authorization and certification procedures
 - Upgrade the curation of the FCC's current databases in preparation for the application of AI/ML
 - Access to AI/ML tools and facilities to perform analysis in support of policies

5.2 Creating an AI/ML Data Task Force Within the FCC

The main recommendation to the FCC related to the data issues is: *The creation of a Task Force to address how the FCC can best address important aspects of data governance and curation to serve its internal needs, and those of industry, and the public - ‘To build knowledge, unleash the Data.’*

One uniformly missing piece in these programs is access to high quality, current, labeled, and curated Data. It’s in the interest of the FCC to help mature the investments in AI/ML relevant to Communications and Operational issues. This cannot be done without the availability of Data and the Sharing of AI/ML Models.

While the research agencies are willing to provide some level of funding for data gathering and collection, by tradition they are unlikely to support the funding for longer term maintenance and curation of the data and the steps that would make the data useful outside the scope of individual programs and projects. This leaves a hole in the larger ecosystem for the contribution that the investments could make to US leadership in Telecommunications.

5.3 Implementing a Procurement Process for Inclusion of AI/ML Capabilities in the FCC's National Broadband Map (NBM)

The AIWG recommends that the FCC develop a holistic approach and plan to satisfy the requirements for the National Broadband Map (NBM). As a step in the process, the FCC should issue an RFI that provides it with the necessary information to choose a procurement approach that best fits the FCC's mandates.

This should include considerations for:

- Identification of technology alternatives, integration approaches, development methods, deployment and operation approaches, and provisions for updating the NBM and updating or replacing the 477 data, as well as estimates of resources needed for each portion of the NBM's life cycle.
- Provisions to use Data gathered for important functions beyond the NBM itself. The purpose is to anticipate how the Data can serve the long term needs of the FCC in developing its AI/ML approaches and does not have to be regenerated (e.g., wireless signal and propagation data).

The development of a Broadband Map has multiple aspects to it and is a major undertaking. It involves capturing the drivers and constraints for the scope, an understanding of and fleshing out of the requirements at a detailed level, operational considerations for how it will be used and by whom, and provisions for eventual operation and sustainment. There are multiple approaches possible and the development, for what it entails, will most likely include a systematic analysis of the trade-offs for how it may be developed, what technologies will be involved, and a concrete plan for how the work will be accomplished (a constructive step by step plan with timelines, milestones, and a budget) and eventually operationalized.

The overall project will also require integration skills, have a significant number of components and elements, field teams, and including data from many different sources.

There is significant capability available commercially and within academic institutions that a well-crafted Request for Information (RFI) will be informative to the FCC. It can help identify technologies available, sources of data, approaches for architecting a solution, approaches for development and eventual deployment, and most importantly providing a factual basis for estimating budgets required for different lifecycle stages.

The results of the response would be followed by an analysis of what role the FCC chooses for itself and what items would be procured and competed for through an RFP.

5.4 Considerations for the Safe Use of AI/ML by the FCC and the Telecommunications Industry

The overarching recommendation to the FCC for the safe use of AI/ML is: *The FCC should establish policies and approaches to ensure the safe use of Artificial Intelligence as it impacts the Nation's Networks, communication needs, and important applications – “Keep humans in control of the loop.”* Stuart Russell (Professor of Computer Science at UC Berkeley) conveyed to our group that while the upside is “an enormous increase in the capabilities of civilization,” there are downsides such as impersonation, bias, robotics replacement of jobs, and the unimaginable outcomes, in cases where AI “outstrips our feeble powers.” Russell suggests AI be designed to be explicitly uncertain about human preferences, so that it turns to humans for judgment (see his slide in Appendix E).

Below we list the finer grained findings and recommendations related to AI/ML safety broken down by the four insertion points discussed in Section 4.4.

(1) Consumer and Enterprise Applications:

Findings:

- Network operators should obey similar guidelines to other providers of applications / content services.⁵
- A future workgroup (with broad representation) should study and report on uses of AI/ML that shape human behaviors.⁶

(2) Operator Business Processes

Findings:

- Operators should obey similar guidelines to other providers of consumer services
- The FCC should encourage the disclosure of information concerning the macro level performance of AI/ML-based services (e.g., demonstrate that customer service, upsell offers, etc. are fair)
- The FCC should consider the merits of a requirement that humans be able to easily determine when they are communicating with a bot (vs. with another human being)

⁵ Our premise here is that operators should not be subject to more stringent guidelines than *over the top providers* and that other regulatory agencies (e.g., the FTC) may weigh in on the use of AI/ML in these contexts.

⁶ Both positive (e.g., economic nudges) and negative (e.g., driving addiction) uses should be considered.

(3) Network Planning

Findings:

- Operators should follow similar guidelines to other providers of critical infrastructure.
- The FCC should encourage operators to disclose information concerning practices used to ensure fairness, e.g., that bias in training sets/models is not a factor in network planning

(4) Network Control

Findings:

- Operators (and their suppliers) should partner with providers of other types of critical infrastructure to adopt and implement best practices with respect to explainability, vulnerability and robustness. These may be similar to existing practices for safety-critical systems.
- It is important for operators to institute processes to manage their AI/ML *supply chain* in order to track the provenance of models and training data used in their networks – and the mechanisms used to secure and manage timely updates to them (e.g., vulnerability patching).
- There is high value in operators sharing, e.g., via an Engineering Consortia, and disclosing the frameworks they have adopted to analyze and address the vulnerability (to attack) and robustness of AI/ML models whose failure could jeopardize the operation of significant portions of the network.

Recommendations for network control and operations:

- The FCC should encourage and incentivize operators (and their suppliers) to institute processes to manage their AI/ML *supply chain*, to track the provenance of models and training data used in their networks – and the mechanisms used to secure and manage timely updates to them (e.g., vulnerability patching).
- The FCC should encourage and incentivize operators to share aspects of the framework they have adopted to analyze/address the vulnerability (to attack) and robustness of AI/ML models whose failure could jeopardize the operation of significant portions of the network.
- The FCC should encourage creation of consortia to develop one or more common assessment frameworks for use with the above analysis and for sharing of critical operational vulnerabilities and lessons learned.

5.5 Initiating Pilot Projects

The AIWG also examined significant and innovative AI projects that focused on important issues such as spectrum sharing, the automated categorization and identification of EM emission sources and the use of AI to improve wireless network protocols and access methods. The work was suggestive of what would be possible in a commercial setting. We perceived a need for a much closer collaboration between the funding agencies, the FCC and Industry, with the objective of steering the research to include conditions likely to be encountered in actual networks.

We also found that Agencies working on Standards, Advanced Prototypes, and Deployments have important capabilities and facilities relevant to the FCC. These include Frameworks for Data Governance, Security and Privacy, etc. Many of these were developed in consortia with Industry. It may be of value for the FCC to be an active partner and participant.

The AIWG held several sessions collecting ideas for where AI/ML may have a significant impact and could be of importance to the FCC's missions. Table 5.5.1 offers a sample of potential areas of exploration for the FCC. A more detailed exposition of selected topics that may result in immediate action for the FCC are detailed in Appendix G.

Area	Benefit
Analysis of Comments to FCC Actions	Better understanding of positions and auto generated material
FCC Data Bases and Website	Improved service for FCC customers and the Public
Network Security and Privacy	Decreased threat exposure
Spectrum Sharing	More dynamic sharing and development of spot markets
Real-time mitigation of Robo-calls	Identification of violations
Emergency Response	Faster service restoration
Preventing Adversarial use of AI	Get ahead of the curve on a rapidly emerging problem - AI Security (AISEC)
Self-Organizing Networks (SON)	Supporting the use of automation for general high-volume applications and critical uses
An Interference Data Exchange	A mechanism to eliminate common causes of interference
AI/ML Based EM Propagation models	Improving the specifications for avoiding interference leading to better utilization of spectrum
AI/ML Benchmarking	Transparency for understanding the performance and behavior of AI/ML models
Emulation of RF Environments	Identification of violations
Automated Testing and Certification Using AI/ML Tools	Dealing with the increasing complexity of software driven devices on the Network
Detection and Elimination of Interference	Improved detection methods and specific identification of sources (Spectrum fingerprinting)

Table 5.5.1: Possible Pilot Projects for the FCC

6 General References

1. AI Working Group Final Report to the FCC TAC, 2019.
2. Crawshaw, J., et. al., “AI in Telecommunication operations: Opportunities and Obstacles,” white paper in Heavy Reading, 2018.
3. “Machine Learning for Communications Emerging Technologies Initiative: Research Library,” IEEE Communication Society, (<https://mlc.committees.comsoc.org/research-library/>).
4. Bootle, R., 2019, “The AI Economy: Work, Wealth and Welfare in the Age of the Robot,” London UK, Nicholas Brealey Books,
5. Gilbert, M., Editor, 2019, “Artificial Intelligence for Autonomous Networks,” Boca Raton FL, Francis and Taylor (CRC Press).
6. Goodfellow, I., Bengio, Y., and Courville, A., 2016, “Deep Learning,” Cambridge MA, MIT Press.
7. Iansiti, M., and Lakhani, K.R., 2020, “Competing in the Age of AI,” Boston MA, Harvard Business Review Press.
8. Kane, G.C., Phillips, A.N., Copulsky, J.R., and Andreus, G.R., 2019, “The Technology Fallacy: How People are the Real Key to Digital Transformation,” Cambridge, MA, MIT Press.
9. Kelleher, J.D., Namee, B.M., and D’Arcy, A., 2015, “Machine Learning for Predictive Data Analytics,” Cambridge MA, MIT Press
10. Knox, S.W., 2018, “Machine Learning: A Concise Introduction,” Hoboken NJ, John Wiley & Sons.
11. Marcus, G., and Davis, E., 2019, “Rebooting AI: Building Artificial Intelligence we can Trust,” New York NY, Penguin Random House.
12. Nawrocki, M., 2019, “Evolution to an AI Enabled Networks,” Presentation to FCC AIWG.
13. Russell, S., 2019, “Human Compatible: Artificial Intelligence and the Problem of Control,” New York NY, Penguin Random House.
14. Saldanha, T., 2019, “Why Digital Transformations Fail,” Oakland CA, Berrett-Koehler Publishers Inc.
15. Siebel, T.M., 2019, “Digital Transformation,” New York NY, Rosetta Books.

16. Westerman, G., Bonnet, D., and McAfee, A., 2015, “Leading Digital,” Boston MA, Harvard Business Review Press.
17. Executive Order 13859, 2019, “Maintaining American Leadership in Artificial Intelligence,” <https://www.whitehouse.gov/presidential-actions/executive-order-maintaining-american-leadership-artificial-intelligence/>
18. Select Committee on Artificial Intelligence, National Science and Technology Council, 2019, “The National Artificial Intelligence Research and Development Strategic Plan: 2019 Update,” NITRD, Washington, DC. <https://www.nitrd.gov/pubs/National-AI-RD-Strategy-2019.pdf>
19. The NIST Big Data Interoperability Framework consists of multiple volumes and was created by NIST in collaboration with over 800 Industrial and Commercial Firms. The framework is described in NIST publication NIST-SP1500 and the individual volumes can be found at: https://bigdatawg.nist.gov/V3_output_docs.php

7 Appendix A. AI/ML Landscape & Annotated Bibliography¹

Introduction

This review of academic literature was undertaken to support the work of the FCC’s Technological Advisory Council on the applications of artificial intelligence in telecommunications. This project’s goal was to help the Technological Advisory Council’s Artificial Intelligence Working Group ensure full coverage of the subject area by analyzing the uses of AI in the telecommunications space, with a specific focus on AI applications in 5G and wireless networks.

As an additional component of the project, research was conducted on general business applications of AI. This area was explored because customer-focused and business-focused AI applications are relevant to the telecommunications industry at large. As a result, issues such as AI in advertising, sales, customer churn prediction, and analytics received attention within this review.

This review of work in the areas of AI in telecommunications and business was performed through the collection of recent, relevant materials published in peer-reviewed journals. This review was conducted through basic Boolean searches performed on the Chinook platform, the University of Colorado Boulder Libraries Catalog. For the purposes of this review, the searches were restricted to journals which published peer-reviewed work, with

Searches Performed for this Review	Number of Articles Returned
(“artificial intelligence” OR AI) AND (telecommunications)	47,318 articles returned
(“artificial Intelligence” OR AI) AND (5G OR “wireless network”)	20,114 articles returned
(“artificial intelligence” OR AI) AND (advertising OR marketing)	45,807 articles returned
(“artificial intelligence” OR AI) AND (churn)	182 articles returned
(“artificial intelligence” OR AI) AND (customer)	7,025 articles returned

Table A.1. Searches performed for this review

¹ This appendix was prepared by Gabriel Lennon (University of Colorado) (Gabriel.Lennon@Colorado.edu). The TAC wishes to express its appreciation for the time and effort Gabriel put into preparing the document.

the returned search results sorted by relevance. The details of the algorithm which prioritizes and returns search were not available, and as such it is worth noting that there is an unknown factor operating in the background with respect to what is returned in these searches. Therefore, it would be beneficial for further work to be done to determine exactly how representative of the total body of work the returned results are. The searches performed in this work were designed to be broad, with the goal of returning general topic treatments as well as research on more specific applications of AI. The searches performed and reviewed are given in **Table A.1**.

To illustrate the volume of content available in this subject area, **Table A.2** includes some proposed additional searches (not included in this review) and the number of returned results.

Example Additional Search Terms	Number of Articles Returned
("Artificial Intelligence" OR AI) AND ("network planning" OR "network management")	584 articles returned
("backpropagation" OR "back propagation") AND (telecommunication)	458 articles returned
("artificial intelligence" OR AI) AND (3GPP)	989 articles returned
("artificial intelligence" OR AI) AND ("telecommunication policy")	312 articles returned
("evolutionary algorithm" OR "genetic algorithm") AND (telecommunication)	12,374 articles returned
("artificial intelligence" OR AI) AND ("vehicular communication" or "vehicular network")	251 results returned

Table A.2. Example of additional searches and the volume of their returned results. These searches were not reviewed for this work

Given the volume of returned results, it would have been impossible to evaluate every returned article within the time allotted for this project, so review was limited to the first 300 articles within each AI-centric Boolean search, and the first 100 articles within each business-centric Boolean search described in **Table A.1**. Supplemental works from January

2019 onward have been taken from NYU WIRELESS, a research center at the Tandon School of Engineering.

Not every result reviewed has been included, as factors such as relevance, coherence, clarity, and overall quality were taken into consideration. This stage of review was subjective. Some of the returned results contained no relevant content, others treated AI as a “magic black box” and the discussed applications of AI appeared too speculative or ungrounded to merit inclusion. Some theoretical and speculative articles have been included nevertheless, as their value to the discussion of AI applications outweighed their purely speculative components. The twenty-three articles selected for this final report were deemed to be of sufficient quality and relevance to merit inclusion. An additional factor considered in this subjective review phase was whether the selected content supplemented the other chosen pieces by providing coverage of additional aspects and areas of AI applications in telecommunications. The goal was to provide wide coverage of the subject area, touching on some of the most researched or most relevant issues in the AI and telecommunications space.

The following review is broken into eight categories consisting of four general telecommunications applications and four general business applications. The review of AI in wireless network applications is organized by the categories of General Treatments, Network Design, Management, and Adaptation, Cognitive Radio Applications, and AI in Vehicular Communications. The review of AI in business applications is organized by the categories of Marketing and Advertising, Churn Prediction, Customer Interaction, and Business Applications of Neuro-fuzzy Systems. The selected articles within each category are addressed individually, with a summary of their contents appearing under the article title. Figures and tables are taken directly from the article under which they are included. A glossary of some of the relevant artificial intelligence terms appears on page 78.

General Treatments

Machine Learning Paradigms for Next-Generation Wireless Networks^[1]

This general treatment on the topic of AI use in telecommunications divides the space of potential machine learning approaches into three parts: supervised learning, unsupervised learning, and reinforcement learning.

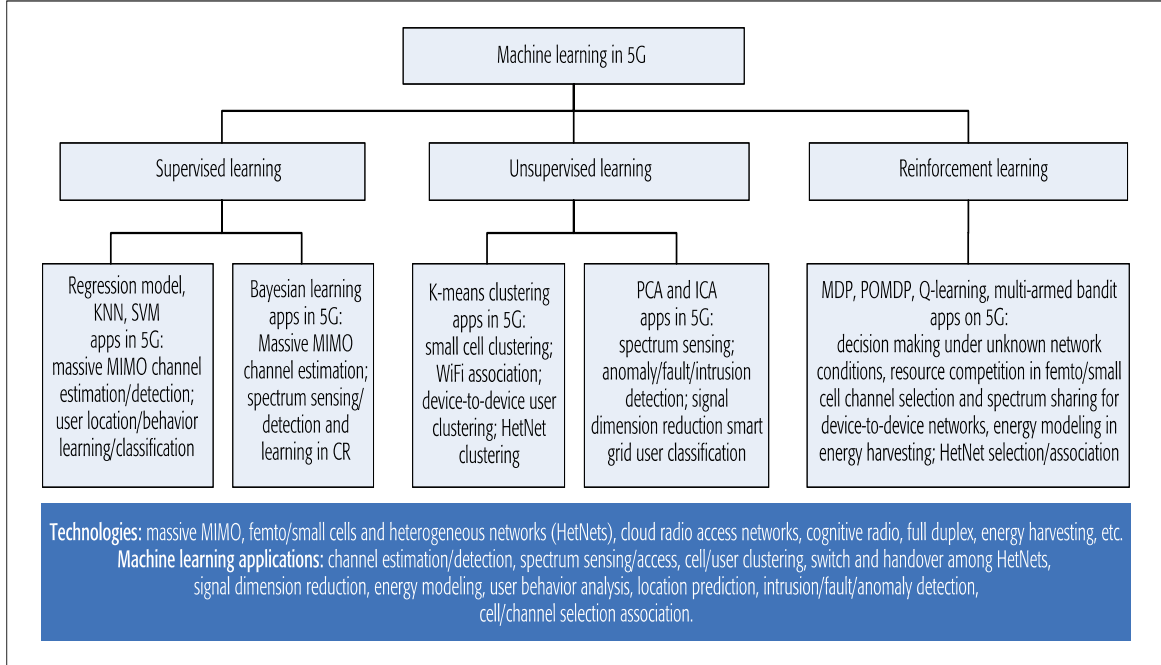


Figure A.1. Machine Learning Approaches in 5G

In addition to the uses specified in **Figure A.1**, supervised learning applications include the use of *K*-nearest Neighbor (KNN) algorithms and Support-vector Machines (SVMs) in heterogeneous networks requiring frequent handovers with the goal of determining an optimal handover solution. The KNN and SVM models can also be used to determine a particular mobile terminal’s usage patterns, dynamically instigating system state changes for the purposes of saving energy without significantly compromising the user experience. This work also identifies experiments which have shown that “up to 90 percent successful energy demand prediction is possible with the aid of the KNN algorithms.”

Like regression, KNN, and SVMs, Bayesian learning can also be productively applied in the massive MIMO arena. In one cited work, the authors addressed the problem of “pilot contamination” by using Bayesian learning methods to estimate the channel parameters of interfering signals and desired links to address a MIMO system’s inaccurate channel estimation. Bayesian learning additionally finds uses in the learning and estimation of spectral characteristics, finding particular utility in cognitive radio networks. Bayesian learning can be applied in this space to determine or estimate primary user behavior, signal strength, inactive states, and so on, combining this information with other variables such as noise variance on sub-bands to generate cooperative spectrum sharing schemes. **Table A.3** provides a concise summary of the applications of supervised learning models within the area of wireless networks.

Category	Learning techniques	Key characteristics	Application in 5G
Supervised learning	Regression models	<ul style="list-style-type: none"> • Estimate the variables' relationships • Linear and logistics regression 	Energy learning [5]
	K-nearest neighbor	<ul style="list-style-type: none"> • Majority vote of neighbors 	Energy learning [5]
	Support vector machines	<ul style="list-style-type: none"> • Non-linear mapping to high dimension • Separate hyperplane classification 	MIMO channel learning [4]
	Bayesian learning	<ul style="list-style-type: none"> • <i>A posteriori</i> distribution calculation • GM, EM, and HMM 	<ul style="list-style-type: none"> • Massive MIMO learning [6] • Cognitive spectrum learning [7–9]

Table A.3. Supervised Learning Models

The authors focus their discussion of unsupervised learning in wireless networks on three specific learning techniques: *K*-means clustering, principal component analysis, and independent component analysis. *K*-means clustering has been demonstrated as a way to efficiently cluster small cells, mobile users, wireless devices, and so on within heterogeneous networks. Principal component analysis and independent component analysis have seen a wide variety of uses as well, as together they naturally lend themselves to signal processing applications such as recovering independent signals from a mixture. Applications such as anomaly, fault, and intrusion detection are also considered within the context of wireless network traffic monitoring. A summary of the uses of unsupervised learning identified by the authors is provided in **Table A.4**.

Category	Learning techniques	Key characteristics	Application in 5G
Unsupervised learning	K-means clustering	<ul style="list-style-type: none"> • K partition clustering • Iterative updating algorithm 	Heterogeneous networks [10]
	PCA	<ul style="list-style-type: none"> • Orthogonal transformation 	Smart grid [11]
	ICA	<ul style="list-style-type: none"> • Reveal hidden independent factors 	Spectrum learning in cognitive radio [12]

Table A.4. Unsupervised Learning Models

Reinforcement learning is put forward as a suitable approach to problems of energy harvesting and resource allocation, as well as more general applications such as cognitive radio networks and channel sensing. The authors describe how the Markov decision process model could be applied to power problems. By considering variables like battery limitations, battery usage, packet transmission/reception rates, and channel selection, the model could guide transmission power in an efficient manner. When used in combination with the Markov decision process, Q-learning (a type of reinforcement learning) may be used as a mechanism for the self-configuration and self-optimization of femtocells. The model described would identify viable unused spectrum and select appropriate sub-channels and

restrictions on which the femtocells could operate while still meeting quality of service standards.

The so-called “multi-armed bandit” model is also proposed as a method of modeling and solving complex wireless resource allocation problems, with the authors identifying the utility of this model as greatest when applied in “multi-player adaptive decision-making problems where selfish players infer an optimal joint action profile from their successive interactions with a dynamic environment, and finally settle at some equilibrium point.” The authors note that this approach is particularly well-suited to device-to-device channel selection problems in cellular networks. A brief summary of the applications and characteristics of reinforcement learning is given in **Table A.5**. The authors conclude their review of the major kinds of machine learning algorithms by noting that there is still a tremendous amount to be learned in the field of AI-enhanced networking.

Category	Learning techniques	Key characteristics	Application in 5G
Reinforcement learning	MDP/POMDP	<ul style="list-style-type: none"> • Bellman equation maximization • Value iteration algorithm 	Energy harvesting [13]
	Q-learning	<ul style="list-style-type: none"> • Unknown system transition model • Q-function maximization 	Femto and small cells [14, 15]
	Multi-armed bandit	<ul style="list-style-type: none"> • Exploration vs. exploitation • Multi-armed bandit game 	Device-to-device networks [16]

Table A.5. Reinforcement Learning Models

Artificial Intelligence for 5G and Beyond 5G: Implementations, Algorithms, and Optimizations^[2]

In this work, the authors identify some of the applications of deep learning approaches within the 5G and “beyond 5G” spaces. The authors highlight ongoing research on AI-based solutions for MIMO detection and multi-user MIMO precoding, as well as AI-based channel coding and precoding with a focus on low-density parity-check codes and polar codes as part of the 5G standardization effort.

The authors describe how, in addition to channel coding/decoding, AI can assist in the creation of other baseband modules such as non-orthogonal multiple access (NOMA) detectors, channel equalizers, and channel estimators. Different neural network configurations are considered with respect to the problem of joint channel decoding and equalization, with the authors advocating for a combination of convolutional, deep, and recurrent neural networks as an adaptive and efficient solution. The authors also provide an overview of how machine learning techniques are appropriate in scenarios where the system model in question is not itself well known. Three examples of this application type are

provided: in-band full-duplex self-interference cancellation, pre-distortion, and the tracking and estimation of fading channels.

In the case of self-interference cancellation, machine learning techniques (namely neural networks) are used to model and reconstruct the non-linear interference introduced into a signal by the “real world” radio frequency components. The use of in-band full-duplex communication has implications for efficient spectrum sharing regimes, and a method of mitigating this self-interference would be a significant improvement in the efficiency of such an approach. The use of neural networks in the performance of pre-distortion to reduce negative impacts on the quality of a transmitted signal is considered, and the authors cite research done in this area indicating that the machine learning approach is at the very least a match for the current conventional approaches to the problem of power amplifier impact on signal quality. Finally, the authors consider the application of machine learning to both estimate and track the behavior of fading wide-band channels. The difficulty of creating a generally applicable model makes the application of machine learning particularly appealing, and the authors cite an example in which a deep neural network was used to predict the behavior of a channel based on a combination of pilots and previous estimates.

The authors also discuss the use of neural networks in “location fingerprinting,” which may be used in scenarios when traditional methods of localization are ineffective or unavailable, and how some components of the fingerprinting process can be used in “channel charting” to enable so-called “self-supervised location sensing.” The authors conclude with the hope that this overview of the general research problems in 5G and beyond 5G will prompt further research in the areas identified.

Network Design, Management, and Adaptation

One area that has received a significant amount of attention within the available literature is the application of AI and machine learning to problems of network design and management, as well as AI-empowered implementations of adaptive networks.

Path Loss Prediction Based on Machine Learning: Principle, Method, and Data Expansion [3]

This article examines three supervised learning-based approaches to the problem of accurately predicting path loss in radio wave propagation: artificial neural networks (also simply called neural networks), support vector regression, and decision trees.

Artificial neural networks are described as a popular approach to solving path loss prediction problems, particularly when provided a large enough sample size. However, they are susceptible to problems of overtraining. ANNs perform well in scenarios where the actual data is a close match for the training data, but they do not fare as well when used with data that is significantly different from the training dataset. Therefore, this method is not well suited to generalization if over-trained (i.e., the model can very accurately predict training examples but struggles to generalize its approach to “real-world” data). Another shortcoming of the neural network approach is that its ability to approximate very complex relationships inherently makes it difficult to explain the predicted results and how the

learning process affected those results. The support vector machine approach (of which support vector regression is an extension) is more able to be generalized than neural networks, with the main challenges being significant computational complexity and the difficulty of defining the functions which transform data into the desired form. Single decision trees run the risk of modeling errors in the form of overfitting, so instead random forests (collections of single decision trees) can be used to predictively model path loss. While decision trees are more easily described and their results more easily understood, this model has shortcomings when it comes to correlating features of the data being used.

To test these three approaches, the authors performed a measurement campaign in an urban microcell in Beijing, China. An omnidirectional antenna was fitted to the roof of a car which drove around the area, attempting to receive signals from a nearby base station. The received signal power was recorded, and the path loss values calculated in a post-processing step which mapped the data to locations within the urban area. The path loss at all positions within the collected data set was predicted with each model, and the results were compared. This comparison revealed that the random forest approach provided the most accurate predictions, followed by support vector regression, neural networks, and finally the traditional “log-distance” model. All machine learning-based models outperformed the log-distance method of predicting path loss.

The authors also address a problem presented in a great deal of the literature: how to acquire enough data on which to train the selected machine learning model. The authors propose two potential solutions: data transferring and combining acquired data sets with classical models. Data transferring is here taken to mean adopting data from other scenarios or other frequencies, provided that the environment in which the data was collected is sufficiently similar. The authors collected data from other base stations in their area and used all these collected samples as training data. The finding was that training data taken from other similar environments and frequencies was sufficient to achieve satisfactory performance of the machine learning models chosen.

The authors also discuss the combination of classical path loss prediction models with machine learning to expand the data set on which the model was trained. This approach is summarized in **Figure A.2**.

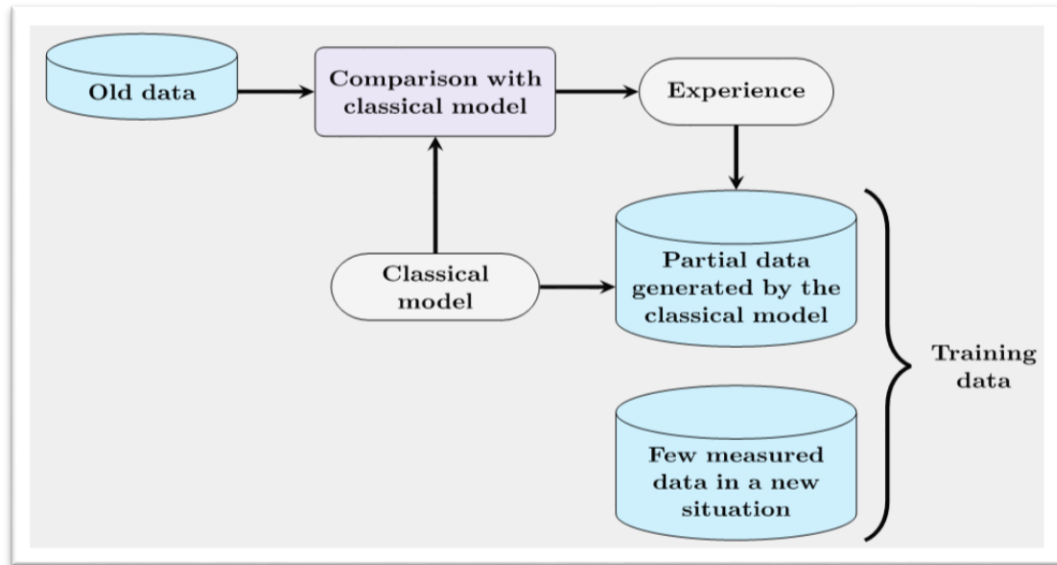


Figure A.2. Expansion of Training Data with the Usage of a Classical Model

To validate this approach to expanding training data sets, the authors utilized an aircraft cabin scenario in which path loss data, signal frequency, and antenna-separation distance were taken from 110 cabin seats. Path loss values estimated by the log-distance model were added to the training data set for the three machine learning techniques, and the results validated the proposal that classical models can be successfully used to expand training data sets. Ultimately, the authors conclude that the three machine learning models which were examined are all in sufficient agreement with measured data to be considered usable for path loss prediction.

Facilitating Mesh Reliability in PPDR Scenarios Utilizing Artificial Intelligence ^[4]

This article examines a novel use of computer vision in the context of public protection and disaster relief. The proposed scheme would utilize computer vision to allow uninterrupted mmWave mesh connectivity, avoiding link blockage. The example given by the authors illustrates a scenario in which disaster relief workers are engaged in fire suppression, and computer vision (here utilizing a convolutional neural network) could be used to examine the immediate surroundings and evaluate the likelihood of a blockage. If link blockage leading to a lost connection was deemed sufficiently likely, additional connections would then be secured (potentially via robotic relays deployed in the immediate vicinity) to avoid any interruption in transmissions. The authors do note that one hurdle to be overcome is that convolutional neural networks are computationally expensive and therefore currently challenging to implement in a real-world surveillance context. The utilization of augmented reality equipment and robotic signal relays described in this paper appears currently more theoretical than practical, however the notion of using computer vision to adaptively reconfigure a mesh network is worth noting.

In this survey, a comprehensive review of the literature on machine learning as applied to self-organizing networks is undertaken. As seen in a previous paper, the authors here identify three main categories of machine learning algorithms to be considered: supervised, unsupervised, and reinforcement learning. The authors categorize each reviewed paper in terms of its proposed solutions and potential use cases and then compare these approaches.

The authors describe three major functions of a self-organizing network: self-configuration, self-optimization, and self-healing. network self-configuration is defined as “the process of automatically configuring all parameters of network equipment, such as BSs, relay stations and femtocells. In addition, self-configuration can also be deployed after the network is already operable. This may happen whenever a new BS is added to the system or if the network is recovering from a fault and needs to reconfigure its parameters.” The authors describe the process of network self-optimization, illustrating how a function must continuously monitor the network (parameters, environment, etc.) and adaptively update the network’s parameters to ensure optimal performance. Self-healing is explained as a function which activates when a fault or failure in the network is detected. The self-healing component of a self-organizing network should not only monitor the network for issues, but it should also be able to diagnose and address them as well.

Given that there are a great many parameters at play in any given network, the complex relationships between these parameters makes the application of machine learning far from trivial. Alteration of one parameter may have an effect elsewhere in the network, fundamentally changing the network’s operation, and capturing all of these intricate relationships within a machine learning-empowered self-optimization model can be challenging.

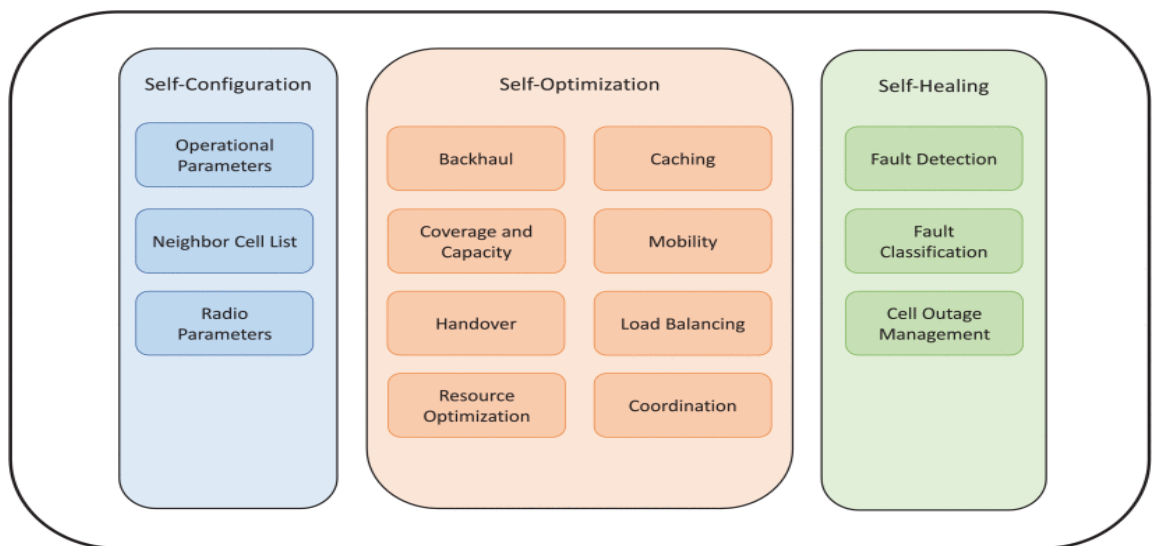


Figure A.3. Major Use Cases of Each Self-organized Network Function

One proposed application of self-optimization is to monitor and evaluate the backhaul connection in a cellular network, however the authors note that despite this being a promising area, there is comparatively little research being performed on this topic. Caching is also considered as a component of network self-optimization, with researchers applying game theory, clustering, transfer learning, and other approaches to attempt to optimize the hit ratio of a cache and therefore provide performance gains within the network.

Finally, the use of machine learning to analyze and optimize network coverage and capacity is considered. Approaches such as optimizing the number of cells within a given cluster in combination with the dynamic adjustment of antenna parameters are considered. The main two sub-problems within this category are identified as antenna parameters and interference control. With respect to antenna parameters, work has been done to optimize traffic offload between macro and microcells, and several authors have addressed the optimization problem by using machine learning to adjust the down-tilt and power of a given antenna to arrive at the optimal parameters.

Several approaches are proposed to address interference control, such as a self-organizing femtocell architecture used to manage the interference between femtocells and macrocells, or the use of feedback controllers to monitor resource usage and assign resources when the resources assigned by the cells themselves are not efficiently allocated. This area has received a considerable amount of attention in the literature, and a great number of possible solutions to the problem of interference control are proposed, many of which have demonstrated promising results.

The authors also identify important aspects of future cellular networks receiving attention in the literature, namely mobility management, handover optimization, load balancing, and resource optimization. Within the world of mobility management, numerous approaches have been proposed to identify a user's current cell and predict the next cell the user will be in. The authors state that current database-based mobility management techniques are inefficient, and that there are large resource gains to be made in this area through the adoption of machine learning solutions. The proposed solutions to mobility management involve the use of machine learning to take in user behavior information and build models predicting any given user's movement pattern. One such solution uses neural networks to predict this behavior by creating mobility models for every user in the network, then making predictions of future behavior based on these generated models. Instead of predicting a user's location, another solution attempts to infer what the user's next activity will be, basing its location prediction off of factors such as time of day, the user's current activity, and the mapping of candidates for the user's next location based on the Google Places API.

The authors identify a significant amount of research on handover optimization, highlighting approaches ranging from probabilistic neural networks to fuzzy logic to Q-learning. Some approaches primarily focus on determining the optimal handover in a given situation, while others attempt to tackle optimal handover as well as load balancing between relevant cells. A chief goal shared between solutions is the reduction of the number of handovers required while maximizing cell capacity at the same time.

With respect to load balancing, the surveyed approaches utilize techniques including regression, feedback controllers, Q-Learning, and genetic algorithms, all applied with the aim of dynamically and optimally balancing variable traffic loads. Research in the general area of resource optimization is also evaluated, with specific focus placed on machine learning approaches to call admission control, energy efficiency, and the coordination of general self-organizing network functions so that they do not interfere with one another.

When addressing the self-healing component of self-organizing cellular networks, the goal is to move away from any need for manual inspection or intervention, instead relying in whole or in part on a machine learning model that can anticipate, find, diagnose, and respond to faults within the network. The research summarized in this survey focuses on three main areas within the self-healing area: fault detection, fault classification, and the management of cells experiencing an outage.

The authors compare some of the algorithms commonly applied in the self-organizing network space on several metrics, such as scalability, training time, accuracy, complexity, and so on. See **Table A.6** and **Table A.7**.

	Scalability	Training Time	Response Time	Training Data	Complexity	Accuracy	Convergence Time	Convergence Reliability
Supervised Learning [53],[77],[260], [261]	Bayes	Low	Low	Low	Low	Fair	-	Fair
	K-NN	Low	Low	Low	Low	Fair	-	Fair
	NN	Fair	High	Low	High	High	-	Fair
	SVM	Fair	Fair	Low	High	High	-	Fair
	DT	High	Low	Low	Low	Low	-	Fair
	CF	High	Fair	Low	High	High	-	High
	K-Means	High	High	Low	Low	Fair	-	Low
	SOM	High	Fair	Low	High	Low	-	Fair
	Game Theory	Fair	-	Fair	-	Fair	-	Fair
	AD	Fair	Fair	Fair	High	Fair	-	Fair
Unsupervised Learning [53], [262]	Feedback	High	-	Fair	-	Low	Fair	Low
	FLC	Low	-	Fair	-	Fair	Fair	Low
	QL	Fair	High	Fair	Low	Fair	High	Fair
	FQL	Low	High	Fair	Fair	Fair	High	Fair
Markov [264], [265]	MC	High	Fair	Low	Fair	Low	-	Fair
	HMM	Fair	Fair	Low	Fair	Fair	-	Fair
Heuristics [227], [229]	Heuristics	Low	-	Fair	-	Fair	-	Fair
	GAs	Low	High	High	High	High	-	High

Table A.6. Analysis of Common ML Techniques in Terms of Self-organizing Network Requirements

	Self-Configuration				Self-Optimization						Self-Healing			
	Operational Parameters	NCL	Radio Parameters	Backhaul	Caching	Coverage and Capacity	Mobility	Handover	Load Balancing	Resource Optim.	Coordination	Fault Detection	Fault Classification	Outage Management
Supervised Learning	Bayes						✓	✓		✓			✓	✓
	K-NN	✓					✓	✓		✓		✓	✓	✓
	NN		✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓
	SVM		✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓
	DT	✓	✓			✓		✓	✓	✓			✓	✓
	CF	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Unsupervised Learning	K-Means	✓				✓	✓	✓	✓	✓		✓		✓
	SOM	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓
	Game Theory	✓	✓		✓			✓	✓	✓				
	AD											✓		✓
Controllers	Feedback	✓	✓			✓		✓	✓	✓	✓			✓
	FLC	✓	✓	✓		✓		✓	✓	✓	✓	✓		✓
Reinforcement Learning	QL	✓	✓	✓	✓	✓		✓	✓	✓				✓
	FQL	✓	✓	✓	✓	✓		✓	✓	✓				✓
Markov	MC					✓	✓	✓	✓	✓	✓	✓	✓	✓
	HMM					✓	✓	✓	✓	✓	✓	✓	✓	✓
Heuristics	Heuristics	✓	✓	✓	✓	✓		✓	✓	✓	✓			✓
	GAs	✓	✓	✓	✓	✓		✓	✓	✓				✓

Table A.7. General Guidelines on the Application of ML Algorithms in Self-organizing Network Functions

The authors make recommendations for the direction of potential future research, highlighting the areas that would benefit from more extensive coverage. Among other areas, the authors suggest further research directed at all three identified major self-organization areas (self-configuration, optimization, and healing). It is suggested that the reviewed research has laid a good foundation for future work, but for these machine learning applications to reach the level of maturity required for real-world implementation, further research is critical.

There is also some discussion around the use of machine learning in 5G. No specific research is presented, the authors opting instead for a general overview of possible applications. These include using machine learning to create additional flexibility in network function virtualization, using machine learning in the physical layer (focusing on applications like automatically choosing the best transmission parameters at a given time), and the use of end-to-end connections which evaluate not only the radio access network but also things like backhaul to pair users with connections which suit their specific needs.

The authors conclude with a statement that for future networks to meet increasingly large and complex demands, the use of intelligent solutions and machine learning will be all but necessary to ensure these demands are met.

Optimal and Scalable Caching for 5G Using Reinforcement Learning of Space-Time Popularities^[6]

In this work, the authors present a method to allow base stations to intelligently prefetch and cache files which are predicted to be popular. The authors advocate for a method in which random user requests are modeled and then processed in a reinforcement learning framework to find an optimal caching strategy. This specifically involves the utilization of a scalable Q-learning algorithm at the base station to track and learn from user file requests.

The authors engage in extensive mathematical analysis for all the schemes and components of their proposed solution, and this analysis is applied to measuring the efficiency of various potential cost functions and caching strategies. The authors subsequently administer similarly rigorous mathematical analysis to validate their proposed scalable Q-learning algorithm caching strategy. In addition to its mathematically demonstrated viability, the proposed approach is also noteworthy for its ability to have its caching parameters easily adjusted in the form of small updates, allowing the scheme even more flexibility. The authors conclude by arguing that the strong results of their mathematical evaluations indicate that their proposed solution is a novel and viable approach to caching problems in 5G networks.

Machine Learning Aided Context-Aware Self-Healing Management for Ultra Dense Networks with QoS Provisions^[7]

In this piece, the authors focus on a self-healing method targeting ultra-dense small cell networks. Algorithms for the detection of and compensation for small cell outages are proposed, both in situations where key performance indicators are available and when they are not. The self-healing scheme developed in this paper is composed of two parts, a small

cell outage detection mechanism (SCOD), and a small cell outage compensation mechanism (SCOC). The SCOD algorithm is designed to take in contextual information in the form of selected key performance indicators and information on the user's position, using this information to detect and locate network misconfigurations as well as sleeping cells. If key performance indicators within the network fall outside optimal bounds, then the algorithm is triggered and the search for the outage begins.

The SCOC algorithm is designed to handle load balancing in these small cell networks. After the detection of an outage by the SCOD, the SCOC scheme tries to optimize the allocation of the available resources, balancing the load within an identified outage area thereby guaranteeing coverage and quality of service requirements.

Mathematical proofs for the proposed algorithms and theorems are provided, and the authors additionally deploy their proposed solution in a simulated environment of 10 small cells and 50 users, finding their approaches to outage detection and compensation to be practical and effective. The authors' overall approach is summarized in **Figure A.4**.

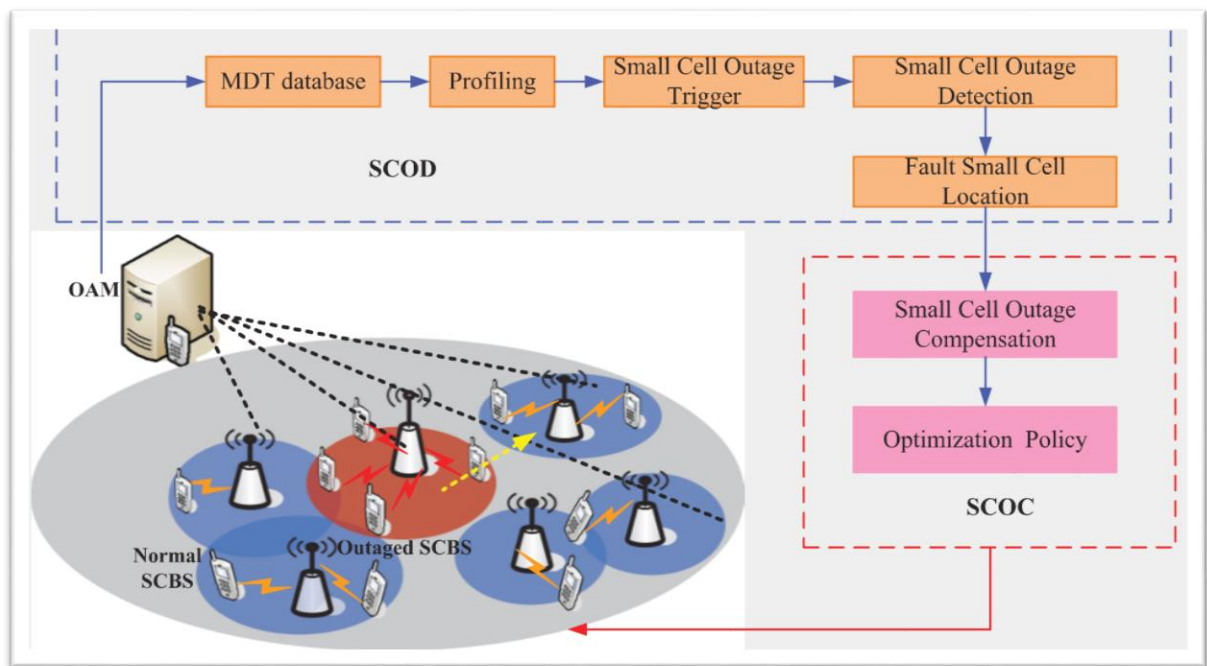


Figure A.4. The Self-healing Framework of Self-organizing Network-based Small Cells

From 4G to 5G: Self-organized Network Management Meets Machine Learning ^[8]

This work identifies the same three foundational components of self-organized networks as the survey above (self-configuration, self-optimization, and self-healing), and discusses the general classes of problems to be addressed within each of these spaces. The main areas identified are variable estimation or classification (such as estimating the QoS of the

network, predicting behaviors to optimize the network, etc.), correct diagnosis of network faults, eliminating the noise in the data generated by the network to allow for accurate processing, pattern identification, and correctly making sequential decisions to adjust network parameters. As in works discussed previously, supervised, unsupervised, and reinforcement learning are discussed in some detail.

Supervised learning is identified as most useful when used to create predictive models of given variables within the network. Six examples of machine learning within the supervised learning category are given, and their specific advantages and shortcomings are briefly discussed.

Unsupervised learning is proposed as a great fit for problems requiring the identification of anomalies within the self-organizing network. It is also well suited to reducing noise within the data and subsequently using that data to identify patterns within the network. The goal of unsupervised learning is described as the ability to predict future inputs without having to feed the algorithm the right answer, instead allowing the algorithm to discover the patterns and solutions on its own. The authors specify that unsupervised learning is best applied within self-organized networks to problems of self-optimization and self-healing, rather than self-configuration.

Reinforcement learning approaches are also discussed, with a particular focus on their use in network management applications which require the adjustment of network parameters themselves. Through “rewarding” or “penalizing” the decision-making part of the reinforcement learning algorithm, it is taught when it is performing well or poorly. This makes reinforcement learning a strong candidate for problems requiring sequences of decisions, as the algorithm will learn from each interaction with the given environment, eventually discovering how to achieve the goal it has been given.

The authors also address leveraging the data generated in mobile networks during the course of everyday operation, advocating for exploiting this information to enable the self-organizing network management scheme to make informed decisions. The suggested data elements are detailed in **Table A.8**.

Information elements relevant for ML enabled SONs.

Source	Data	Usage
Charging Data Records (CDR)	Includes statistics at the service, bearer and IP Multimedia System (IMS) levels.	These records are typically stored, but only used by customer service. The network operation departments typically do not leverage this information and do not have access to it, as much as customer service does not leverage network management data.
Performance management (data on network performance)	It covers long-term network operation functionalities, such as Fault, Configuration, Accounting, Performance and Security management (FCAPS), as well as customer and terminal management. An example is that defined for OAM, which consists of aggregated statistics on network performance, such as number of active users, active bearers, successful/failed handover events, etc. per BS, as well as information gathered by means of active probing.	The data is currently mostly used for fault identification, e.g., triggering alarms when some performance indicator passes some threshold, so that an engineer can investigate and fix the problem. Typically, the only automatic use of this info is threshold-based triggering, which can be done with very low computational complexity.
Minimization of Drive Tests (MDT)	Radio measurements for coverage, capacity, mobility optimization, QoS optimization/verification	This data is used for identified use cases such as coverage, mobility and capacity optimization, and QoS verification
E-UTRA Control plane protocols and interfaces	Control information related to regular short-term network operation, covering functionalities such as call/session set-up, release and maintenance, security, QoS, idle and connected mode mobility, and radio resource control.	A This information is normally discarded after network operation purposes have been fulfilled. Some data can be gathered via tracing functionality or used by SON algorithms which normally discards the information after usage

Table A.8. Information Elements Relevant for ML-enabled Self-organizing Networks

The authors also review some of the relevant literature on machine learning in network management. The topics touched on are mobility, load balancing and optimization, coverage and capacity optimization, the coordination of inter-cell interference, energy savings, cell outage detection and compensation, coordination of conflicting self-organizing network functions, minimizing drive tests, core networks, and virtualized and “softwareized” 5G architectures.

The authors conclude by stating that machine learning is a “crucial and inevitable” tool to address the current and expanding need for network automation. The authors suggest that the huge amount of data being generated in current cellular networks is being underutilized with respect to self-organized networking and argue that self-organized networking should be adopted as a core component of 5G network management.

Multi-layered Intrusion Detection and Prevention in the SDN/NFV Enabled Cloud of 5G Networks using AI-based Defense Mechanisms ^[9]

The authors here investigate the potential of a machine learning-based intrusion detection and prevention system within the network function virtualization and software defined networking-enabled “cloud” of 5G networks. The particular attacks focused on by the authors are IP spoofing, flow table overloading, DDoS, control plane saturation, and host location hijacking.

A so-called “multilayered intrusion detection system” or “ML-IDP” is proposed, and the authors divide the security issues it seeks to address into five layers: the data acquisition

layer, switches layer, domain controllers layer, smart controller layer, and virtualization layer.

Techniques like robust user authentication, intelligent packet classification, and the selection of optimal switches for an incoming request are examined and applied in a network simulation. The results of the simulation are evaluated through several metrics such as switch failure rate, the detection rate of malicious packets, and packet delay from source to destination. Comparative analysis is performed in which the proposed intrusion detection system is compared against four other attack mitigation methods. In conclusion the authors argue that their AI-based approach is superior to the alternatives regarding the specific threats examined.

Real Time Scheduling and Power Allocation using Deep Neural Networks^[10]

In this work, deep neural networks and deep Q-learning are applied to estimate optimal power control and link scheduling in a scenario with a cluster of multiple interfering small cells. The authors' novel approach is justified by the high computational time cost of another near optimal scheme (a combination of exhaustive search and geometric programming), which makes that solution essentially unusable in a real-time environment. The authors aim to replace this unwieldy solution with their own.

The challenge of efficient power allocation is approached through the application of a deep neural network, designed to allocate power to user equipment and base stations with the goal of maximizing the weighted sum link rates of the system.

The authors next propose an implementation of deep Q-learning called a deep Q-network (DQN) which would involve training the DQN with the data that is generated from the implementation of the power allocation neural network. The system is designed to output viable link schedules as well as power control schemes.

The authors simulate an environment consisting of four randomly distributed small cells, and implement their proposed solution finding that without the power allocation neural network, training the deep Q-network can be too time consuming. However, when the two are correctly combined they can achieve a weighted sum link rate that equals or outperforms the exhaustive search and geometric programming methods. The proposed deep learning method can be implemented to show a reduction of computation time of five orders of magnitude, with less than a nine percent reduction in performance.

The authors state that further work on ways to train the deep neural network and the Q-learning network together would be beneficial, and advocate for the adoption of their proposed system as a replacement for exhaustive search-based methods for scheduling and power allocation.

Distributed Learning Algorithms for Spectrum Sharing in Spatial Random Access Wireless Networks^[11]

Here the authors identify two main spectrum access protocol "classes": those in which users share information with one another and those in which users do not. The authors state that

to effectively manage spectrum access, the users must adapt their actions based on provided information about the system they are in. They argue that stable channel allocation to users may not be the optimal solution when examining the entire system and investigate whether there is a way to create an optimal spectrum sharing solution for the whole system based only on small amounts of shared information.

For cases where users are not cooperatively sharing information, the authors propose a distributed learning approach with each user taking a selfish approach by trying to maximize their own rate without consideration of other users. In the case where users *are* cooperatively sharing information, the authors advocate for a different distributed learning model which would function by attempting to achieve what they refer to as “global proportional fairness” with regard to spectrum sharing within the system.

The authors run their models through simulations and provide extensive mathematical proofs of the functionality of their designs, concluding that their algorithms work well within the problem space. However, the authors note that the computation of the proposed optimal solution may not be possible from a practical standpoint depending on the size of the network in question.

Cognitive Radio Applications

Artificial Intelligence-based Cognitive Routing for Cognitive Radio Networks^[12]

The authors of this survey describe what they feel has been a gap in the work done on AI in cognitive radio networks, namely that most efforts have focused on optimization at the level of the individual cognitive radio nodes, without sufficient focus on network-wide optimization. The authors advocate for the creation of cognitive routing protocols which would be incorporated into the design of future routing protocols for cognitive radio networks.

The authors acknowledge that spectrum awareness is certainly included in earlier cognitive radio network routing protocols but claim that in the future AI-based techniques should be included in network design and architecture. Several learning techniques are described and evaluated by the authors, and the strengths and weaknesses of each are addressed. The categories of supervised, unsupervised, and reinforcement learning are also discussed, and their applications and challenges are considered.

The authors discuss supervised learning models and conclude that given the need for these models to be trained on historical data, they are not particularly suitable for the routing applications being examined. Supervised learning has seen significant use in the world of cognitive radio networks, with the authors identifying its primary use as signal classification (which in the context of routing is largely only relevant when considering adjacent problems such as primary user detection). In supervised learning, a major assumption commonly made is that the environment in which the model will operate is somewhat static or “stationary,” and therefore the training data will also be so. The authors highlight the potential pitfall of “concept drift,” an issue in which the target variable the model is attempting to predict changes over time in unanticipated ways. This is a danger to be

cognizant of when attempting to use supervised learning within this context. Since the environment being interacted with in cognitive radio networks is typically “non-stationary,” supervised learning models are expected to play at most a minor role in the potentially unknown RF environments that are being discussed.

The objective of unsupervised learning here is to be able to create a model of the provided inputs without being given any information on what the corresponding outputs should be (in contrast to supervised learning). The authors state that unsupervised learning can be readily applied to problems in which it is necessary to extract information about signals based on particular measurements (useful for applications such as identifying and classifying how unidentified cognitive radio network systems are using spectrum). Regarding its application in routing, the authors state that unsupervised learning is well-suited, as it does not require training data or supervision to address relevant routing problems. That said, the authors identify overfitting and a long convergence time as potential drawbacks to this approach.

Reinforcement learning is discussed in combination with Markov Decision Processes (MDPs), a decision-making framework particularly well suited to use with reinforcement learning. Reinforcement learning is cast as a collection of “learning problems” rather than a prescribed set of algorithms or approaches. The authors discuss a number of reinforcement learning methods, beginning with Q-learning. Q-learning is described as very useful within routing problems, and the evolution of this approach to routing is detailed. Q-learning has been used in routing methods which are designed to minimize delivery time by using the model to experiment with and learn from diverse routing policies. Most applications of Q-learning in this context involve leveraging its ability to “explore” a number of different strategies, learning from each new attempt.

The authors propose that because cognitive radios are often deployed to work in “unknown” environments, reinforcement learning models are particularly well-suited to meeting this challenge. The authors cite other works which have explored the general topic of reinforcement learning in cognitive radio and discuss the numerous ways in which reinforcement learning can be applied to the specific problem of routing. The authors highlight the use of reinforcement learning in more general routing schemes and note that its main drawback in the world of cognitive radio networks is slow convergence. The authors caution that while reinforcement learning is a good candidate for routing in cognitive radio networks, the necessity of learning complex network dynamics necessitates a similarly complex reinforcement learning implementation. In order for it to operate effectively in a constantly changing network environment, a great deal of care must go into the construction of the reinforcement learning-based solution.

The authors next explain online learning algorithms, a form of machine learning that operates by making sequential decisions given only partial information. The online “multi armed bandit” scheme here discussed has been applied to many cognitive radio network problems, such as management of opportunistic spectrum access. The authors note that the dynamic environment of cognitive radio networks makes the application of the so-called

“restless bandit” framework particularly suitable. While this scheme has been used in the literature in spectrum sensing applications, there has been no previous work done on it with respect to routing in cognitive radio networks. The authors comment that one of the shortcomings of the multi-armed bandit approach to routing is that since it is considered a “naïve” approach, the framework can yield poor performance for a number of reasons including that multi-armed bandits cannot model the type of multi-user interaction common in cognitive radio networks.

Game theory is the next framework examined by the authors, and they note that a combination of game theory and the ability to “learn” from past experience has received a lot of attention in the prior literature. One cited work detailed how to use this “learning with game theory” approach to address resource management in multi-hop cognitive radio networks, and another discussed how to use a similar approach in heterogeneous 4G networks. However, the authors of the present work note that while this area appears promising, limited work has been done. One of the main concerns with using this approach in cognitive radio network routing is that the information available to the framework is limited or incomplete, and it can be challenging to gain a holistic picture of the “players” as well as the relevant strategies and network environment.

The authors next turn to genetic algorithms, a type of evolutionary algorithm. The authors describe genetic algorithms as “a very general meta-heuristic technique which can be thought of as the sledgehammer of the craft of algorithms.” The genetic algorithm approach includes the use of recombination and random “mutations” paired with subsequent evaluation to improve on the current “best” solution. Within the context of cognitive radio networks, the authors state that genetic algorithms have been “extensively deployed” and have seen use not only in routing applications but in problems such as bandwidth allocation as well. The use of genetic algorithms to address routing in cognitive radio networks includes the non-trivial task of defining an evaluation function for the results, in addition to the potential drawback of slow convergence. The authors note that to some extent these challenges can be mitigated by identifying a way to terminate the evolution when the learned solution is “good enough.”

So-called “ant colony optimization” (another form of evolutionary algorithm) is discussed next, which can be used as a method for approaching “shortest path” problems with limited resources. The authors cite research which has shown that ant colony optimization can outperform genetic algorithms in routing applications, particularly in dynamic environments. However, the authors note that there are no guarantees of good convergence times with regard to a globally optimal solution, making this approach difficult to implement for larger problems.

The literature on artificial neural networks is discussed next, and the authors note that these networks are typically used in a supervised learning context. Despite this, there has been research done which directly applies artificial neural networks to routing problems with notable success. However, within the context of cognitive radio networks, the supervised

nature of artificial neural networks significantly limits the usability of this approach, and the danger of overfitting the solution adds additional difficulty as well.

Bayesian learning (a model designed to make predictions based on the probability of hypotheses) has been proposed for future work, but little research has been done on the use of this model in the context of cognitive radio networks. Bayesian models have been proposed to estimate the activity of primary users and to model spectrum access within cognitive radio networks, but their utility in the context of cognitive radio networks is largely limited to a supportive role.

The authors address the future of cognitive routing protocols, suggesting that they would benefit from the implementation of several prediction-based tasks, namely the modeling of channel quality and spectrum occupancy. The information gained from modeling primary user activity and patterns may be used to aid predictions in learning-based solutions in order to select more stable routes for the given signals.

Finally, the open questions in this area and suggested future work are given. Special attention is paid to the suggestion of additional research on “multi-agent decision making” instead of merely focusing research on “centrally controlled optimization.” Brief attention is also given to various machine learning techniques not otherwise covered in this work, and the combination of machine learning with modern technologies such as software defined networks is briefly addressed. The authors conclude by stating that combining learning from past and present conditions within a cognitive radio network could lead to notable performance gains. **Table A.9** and **Table A.10** contain useful summaries of the approaches discussed in this survey.

Learning Techniques	Applications to Routing	General Applications to CRNs
Hidden Markov model	Can indirectly utilize spectrum occupancy and channel quality predictions	Spectrum occupancy prediction: Akbar and Tranter (2007), Park et al. (2007) and Choi and Hossain (2013); spectrum sensing, primary signal detection (see references in He et al. 2010)
Reinforcement learning	Q-routing algorithm (Boyan and Littman 1994); learning automata (Akbari Torkestani and Meybodi 2010a, b); RL-based routing for CRNs (Xia et al. 2009), MANETs (Bhorkar et al. 2012; Di Felice et al. 2010), see further references in survey papers (Bkassiny et al. 2013; Al-Rawi et al. 2013; Yau et al. 2010)	Dynamic channel selection and topology management (Yau et al. 2010); spectrum sensing and efficient spectrum utilization through PU detection (Reddy 2008); security (He et al. 2010)
Learning with games	Evolutionary game theory; dynamic Bayesian games (Pavlidou and Koltsidas 2008); congestion games (Pavlidou and Koltsidas 2008); quality of routing games (Busch et al. 2012)	Spectrum access games (Van der Schaar and Fu 2009)
Online learning	No regret routing for adhoc networks (Avramopoulos et al. 2008; Bhorkar and Javidi 2010); Bandit routing (Tehrani and Zhao 2013); online adaptive routing (Awerbuch and Kleinberg 2008)	Opportunistic spectrum access (Han et al. 2007)
Genetic algorithms	Shortest path routing (Ahn and Ramakrishna 2002)	Modeling wireless channel: (Rondeau et al. 2004b)
Ant colony optimization	Routing with ACO in CRNs (Zhao et al. 2012) and MANETs (Caro et al. 2005)	Cognitive engine design (Zhao et al. 2012)
Artificial neural networks	Routing with ANNs Ju and Evans (2010) and Barbancho et al. (2006)	Spectrum occupancy prediction (Tumuluru et al. 2010); dynamic channel selection (Baldo et al. 2009); radio parameter adaptation (see ref. in He et al. 2010)
Bayesian learning	Bayesian routing in delay-tolerant-networks (Ahmed and Kanhere 2010)	Establishing PU's activity pattern (Saad et al. 2012; Han et al. 2011); channel estimation (Haykin 2005); channel quality prediction (Xing et al. 2013)

Table A.9. Summary of the Various Learning Techniques Discussed

Technique	Class	Pros	Cons
Optimization techniques	Single-player optimization	Well-developed theory that has been extensively applied. Many interesting practical problems can be formulated as convex optimization problem whose efficient solution methods exist	Requires complete knowledge of the environment which can be impractical in many situations; The theory does not incorporate multiple decision makers or changing environment
Markov decision processes (MDP)	Single-player optimization (with controllable states)	MDP is a useful tool for sequential planning, or control, of dynamic CRN processes. It can model and find optimal actions for situations where the outcome does not follow deterministically from actions	MDPs make two assumptions (the Markovian assumption and assumption of a stationary environment) that may not be realistic in CRNs. MDPs also assume a single player and cannot model the presence of multiple users in CRNs
Hidden Markov models (HMM)	Single-player optimization (with hidden states)	HMM are excellent models of temporal processes and can be used to model and analyze CRN processes such as PU arrivals	The Markovian assumption, fundamental to HMM, is often not satisfied by temporal processes with memory. Also, HMM is a supervised learning technique that is entirely suited for routing optimization
Reinforcement learning (RL)	Unsupervised learning	Can be applied in unknown environments. Can optimally solve Markov decision processes (MDPs)	The main drawback of the RL technique is its slow convergence. RL techniques may not converge towards optimal solution. More efficient methods exist in supervised settings/or when the environment is known
Genetic algorithms (GA)	Metaheuristic optimization	Excellent for parameter optimization and search of solutions for complex problems for which optimal solutions are unavailable/too expensive	The main problem of GAs is slow convergence. Can suffer from local minimas. Genetic algorithms are not well suited to real-time applications
Ant colony optimization (ACO)	Metaheuristic optimization	ACO is an efficient simple metaheuristic that is useful for routing in dynamically changing environments (such as those present in CRNs)	The disadvantages of ACO stem from the heuristic nature of ACO. There is no guarantee that ACO will converge towards an optimal solution, and the time of convergence may be high for large problems
Artificial neural networks (ANNs)	Supervised learning/unsupervised learning	Excellent for classification; does not require prior knowledge of the distribution of observed process; applies very generally to a wide variety of problems; easy to scale; can identify new patterns	Training of ANNs can be slow and can lead to overfitting; No underlying theory to link application with required network
Game theory	Multi-player optimization	Can model interactive optimal decision making between multiple decision makers; can be used to study competition	Requires complete knowledge of the environment which can be impractical in many situations; Convergence to optimal strategy cannot be guaranteed for many games
Bayesian non-parametric (BNP) methods	Supervised learning/unsupervised learning	BNP techniques are primarily tools for inferencing, clustering, planning and prediction. BNP techniques have been widely applied in this context in CRNs	While BNP techniques are useful for a variety of tasks (such as learning and predicting PU activity), their direct application for the direct algorithmic problem of routing in CRNs is limited

Table A.10. Comparison of the Main AI Techniques Presented in this Appendix

Primary User Characterization for Cognitive Radio Wireless Networks using Neural System Based on Deep Learning^[13]

In this work, the authors examine the application of a “Long Short-Term Memory” (LSTM) recurrent neural network to estimate when primary users will leave assigned spectrum vacant so non-licensed users can make use of it. The Long Short-Term Memory neural network is designed to tackle sequence problems, rather than only being able to work on single data points.

The authors first discuss the literature on the topic of algorithmic characterization of primary users in a network, examining the work done with learning-based algorithms, queueing theory-based algorithms, and other methodologies. The mathematics necessary to define the components of the LSTM framework are given, and the authors describe their testing approach in two test cases. In test case one, the authors analyzed the modeling and prediction capabilities of the LSTM neural network when given computer-generated input signals mimicking uncommon behavior patterns. In the second test case, the authors fed the LSTM network real data sequences, and studied the predictions made. They note that in the second use case, the study was done under the assumption that 70% of the data was used in the training stage and 30% was used for validation.

In evaluating the results from test case one, the authors state that LSTM was able to adequately generalize primary user behavior in the cases considered, but they also note that this is somewhat contingent upon the primary user in question continuing to exhibit the same behavior. Despite this shortcoming, the authors also note that the validation error for this approach is very small, which indicates that it is possible to optimally model the network using this approach.

Processing time is higher in the second test case because of the longer traces performed by the LSTM network, as well as its need to store information to be used in the future. The average prediction success in the LSTM approach was 87.34% for GSM and 76.30 for WiFi. The authors propose that the reason the LSTM framework performs better in the GSM spectrum band is likely that WiFi traffic flow tends to be very chaotic in nature. In conclusion, the authors evaluate the performance of the LSTM network approach and find that it is suitable for estimating spectrum usage by primary users.

AI in Vehicular Communications

An Online Context-Aware Machine Learning Algorithm for 5G mmWave Vehicular Communications^[14]

The authors of this paper model environment-aware beam selection in mmWave vehicle systems as a “multi-armed bandit” problem and propose an online learning algorithm which they call “fast machine learning.” The idea is that the algorithm will “explore” different beams while also accounting for factors such as the target vehicle’s direction of travel. This process allows adaptation to system dynamics (like changes in traffic patterns or real-world blockages like buildings) by examining aggregated information from the vehicles

themselves for each beam. The algorithm evaluates this contextual information and then uses it to “re-explore” and adapt the beam selection in response to these changes.

The authors argue that this approach is flexible and generic enough to be adapted to new 5G use cases as needed. The mathematical formulation of the beam selection problem is given from the perspective of the base station, the algorithm for addressing this problem is defined, and a proof of the relevant theorem is cited.

The authors next turn to the feasibility of their proposed “fast machine learning” model and describe its potential application within 4G and 5G systems. This model is broken down into four steps: a registration request from the vehicle, a response from the base station about the selected beams, standard cellular attachment by the vehicle when it enters coverage, and feedback if the vehicle fails to detect the base station within the beam.

Extensive numerical evaluation of the proposed scheme is done via a simulation setting chosen according to 3GPP specifications, and the authors conclude that the “fast machine learning” approach is superior to the benchmarking algorithms to which it was compared. The authors summarize relevant related work and conclude that their simulation results demonstrate the capabilities of “fast machine learning,” including its suitability for use in other areas of mmWave networking.

Marketing and Advertising

The Impact of AI on the Advertising Process: The Chinese Experience^[15]

In this piece, the authors provide observations regarding the impact of artificial intelligence on the marketing process within select Chinese markets and propose a simple framework for conceptualizing the AI-empowered advertising process.

A brief literature review is conducted, focusing mainly on the area of “programmatic buying” (the use of algorithms and software to purchase digital advertising), as well as how other authors have described the changes AI may bring to other facets of advertising such as altering advertising research and impact evaluation processes. The main applications of AI in advertising identified by the authors are research and market analysis, ad design and copywriting, media planning and buying, and performance evaluation and feedback.

In the context of research and market analysis, the authors describe the use of AI to power “consumer insight discovery.” User behavior and patterns are extensively mined, and game theory and AI can be employed as ways to predict user behavior. Information such as GIS and GPS data, as well as social media post content is utilized as well. As an example, they describe how multinational media company Dentsu Aegis Network has used identification and analysis techniques on pictures posted to social networking sites such as Weibo.com to determine why restaurants owned by a client were underperforming in a specific region. The dietary preferences in this province were evaluated at different times and this information was used to generate a clearer picture of why sales had slowed down. The authors note that such an application essentially combines research and analysis into one step.

Regarding ad creation, the authors discuss the commerce website Alibaba's use of AI technologies in ad design. Initially only designing ads for its own platforms, eventually Alibaba began offering its AI-enabled design service to third parties. The authors claim that after allowing others to purchase the service, Alibaba helped design 6 million posters for vendors in a relatively short time. The authors also note that AI-based ad creation competitors have emerged recently. In the world of copywriting, another Chinese commerce platform called JD.com (a main competitor to Alibaba) launched a "smart copywriting" system designed to automatically generate, among other things, product descriptions for vendors on their site. The system they implemented analyzes user searches and quickly generates a matching merchandise description. This description is then saved on a recommendation list for the next time a similar user search is seen.

The authors identify the goal of media planning and buying as largely revolving around the comparison of consumers' "digital lifestyles" with their actual lifestyles in an attempt to optimize the media planned and bought with that consumer in mind, such as personalized ad content.

Evaluating the impact of advertising is done through real-time monitoring of consumer activity, with machine learning models used to create consistency between the impact of the brand and the impact of the advertising. Real-time feedback data is considered, and machine learning models can be employed to meld feedback with observed results. Adjustments to the advertising targets and content are made based on the results of this analysis with the goal of optimizing the responses on the media planning and buying side.

Next the authors turn to a summary of the general characteristics of AI-based advertising, identifying data-based and tool-based approaches, as well as their ability to produce synchronized, highly efficient advertising content. Regarding the data-based approach, the authors claim that the use of AI in advertising allows what was previously a process composed of sequential steps to be condensed into a more parallel approach, allowing the steps discussed above to operate in a "synchronized" manner. The authors claim that this approach shifts the work from a labor force problem to a problem best approached from the world of data and algorithms. In a tool-based approach, software is used to tackle the data collected in the data-based approach, leading to things discussed previously such as targeted, automatically generated ad copy. Efficiency is of course a large consideration, and the authors point out that techniques such as programmatic buying have had tremendous impacts on advertising efficiency, and that AI-based advertising design has reduced costs, increased speed, and lead to an easier method of handling user preferences and ad impact. The authors conclude that while AI-based advertising is causing a dramatic shift in the Chinese advertising world, it has not yet lead to a "reengineered" advertising process. Instead, AI tools and data analysis are enhancing the "traditional" approach composed of the four advertising steps discussed above.

Marketing and Artificial Intelligence^[16]

In this analysis of artificial intelligence applications in the marketing world, the authors attempt to identify how deeply AI has penetrated into this field, examining possible

implications for professionals in the marketing space. A brief and very general discussion of AI techniques such as machine learning and natural language processing is provided before the authors move onto their two primary research questions: do all forms of AI have potential marketing applications (an ambitious question), and what are the impacts that AI-enabled marketing could have on the field itself. To address their first question, the authors provide an extensive table of 5 AI applications and their myriad uses within marketing (see **Table A.11**).

AI areas	Examples of application in marketing
Voice processing technologies	<ul style="list-style-type: none"> • Voice purchase requests made through a device or the Amazon Alexa app. • Virtual assistants are supporting task execution (Siri, Google Home, Cortana).
Text processing technologies	<ul style="list-style-type: none"> • Use of a virtual assistant as a guide to walk you through a shopping centre (Alpine.AI). • A virtual assistant embedded in a mobile bank app, taking advantage of NLP, handles client requests alone by responding to their inquiries. A virtual assistant is presenting application features, options to make a purchase of bank products by oneself, and providing information about the location of bank branches and cash machines (ING Bank Śląski). • A GPS navigation system that apart from showing the route to the

Table A.11. Examples of the Application of AI in Marketing (Continued Below)

	<p>selected destination suggests attractions found nearby or on the way to the destination, and shows similar objects to those related to the set destination (Naver).</p> <ul style="list-style-type: none"> • An analysis of statements made by clients of banks, insurance companies, and telecoms performed in order to diagnose irritating situations, which led to the elimination of negative events that might occur in the customer journey, and to modification of the customer service process (Touchpoint). • Development and launch of new beer recipes, and modification of the existing products thanks to information gathered by a chatbot (Intelligentx Brew). • Development of a marketing campaign to launch a new car model - the Toyota Mirai. Using data provided by a selected target group, computers performed an analysis of texts and videos on YouTube in order to teach the machines the preferred style of the said target group. Next, through multiple iterations, they developed the first creative advertising campaign, and the final texts for the adverts were approved by the supervising team. The result was almost a thousand of advertising spots tailored to the profiles of the ad recipients on Facebook (Toyota, Saatchi&Saatchi). • Promotion of the <i>Milonerzy</i> TV show, the Polish edition of <i>Who Wants to Be a Millionaire?</i>, on Facebook taking advantage of a conversational chatbot. Maintaining the format and the style typical of the show made it possible to offer new and unique experience (TVN).
Image recognition and processing technology	<ul style="list-style-type: none"> • Face recognition as a way to make payments (KFC). • Recognising the condition of face skin, followed by an individual selection of the type of face cream based on an analysis of one's photo and data, including information about the current weather (Shiseido). • A photo as a medium to search for items online. Apart from search results in the form of identical items, the search engine offers similar or complementary items (eBay). • Using the client's face image to select colour cosmetics individually during online shopping (Estée Lauder). • Service-free bricks and mortar shop where video cameras analyse the selected products and payments are made automatically (Amazon). • Electronic mirrors in a clothing shop that match the collection to the client's appearance, style, and taste (FashionAI). • Selection of the best Christmas gift by going through twelve best suggestions. Based on the recognition of the buyer's face and emotion analysis, the programme suggested the best option to go for (eBay). • Identification of clients before the start of a video consultation by comparing the video image with a photo provided earlier by the client (BBVA). • Embedded ML mechanisms make it possible to automatically frame images according to the requirements of the brand and communication channels (Adobe Sensei). • An image finder that makes it possible to select the best photos and reject the less appealing ones (Everypixel).
Decision-making	<ul style="list-style-type: none"> • Development of individual savings plan thanks to an analysis of the funds available on one's account, receipts, amount of expenses and the way one spends their money. By comparing the financial behaviour of a user and a given community, the application develops a tailor-made savings plan to match the financial capabilities of a given person (Plum). • Travel destinations matched individually based on the traveller's musical preferences. Apart from the city, the app chooses specific districts and attractions to match the user's profile (Spotify, Emirates).

	<ul style="list-style-type: none"> • A chatbot is preparing a cocktail recipe using the ingredients the consumer has at home and based on the consumer's preferences. The chatbot analyses 300 recipes and offers the best-matched solution (Diageo Simi Bartender). • Based on the user's mobile phone data (location, sun exposure time), the app indicates the right level of UV protection filter (Monteloeder). • Dynamic matching of prices to the user based on their shopping record visited websites, or the owned mobile phone (iperfummy.pl, kontigo.pl). • Matching adverts to user characteristics based on one's online history (ING Bank Śląski). • New product recommendations (Amazon, Netflix). • GO-I-PACE, an application is analysing one's driving style, route choices, and frequency of charging the car (electric car). Based on the results, the app offers suggestions on how to drive the care in a more efficient and effective manner (Jaguar I-PACE). • ZozoSuit helps customers order clothes fitted perfectly to their figure. Thanks to in-built 150 sensors, ZozoSuit makes it possible to take 150,000 measurements (Start Today, StretchSense). • A platform to manage marketing campaigns online. In the first weeks, AI learns the specificity of a given company, then, based on data analysis, comes up with recommendations concerning the campaign strategy (Albert AI, Harley Davidson). • Detecting faults and errors in product functioning and forecasting malfunction occurrences. The synchronisation of the work performed by the technical team responsible for device (lift) monitoring and repair works (if necessary) (KONE, IBM Watson IoT, Salesforce Einstein). • Creation of a consolidated customer record regardless of the products purchased and used, linking customer data from every company area (Sales Cloud Einstein, U.S. Bank). • Synchronisation of customer data from all possible points of contact with the brand (social media, website, e-mail, phone conversation). All interactions are aggregated and presented in one place in order to offer improved customer service (Salesforce, Adidas).
Autonomous robots and vehicles	<ul style="list-style-type: none"> • Service-free shops (Ford & Alibaba, Amazon Go, Zaitt Brasil). • A robot used to check the stock on shop shelves and the arrangement of the products displayed. Information of shortages or incorrect arrangement is sent to the service staff, who take their time to look into the reported issues (Schnuck). • An autonomous shop is offering basic and fresh products and magazines, able to travel independently to the warehouse in order to replenish the stock. The shop was tested in Shanghai (Moby Mart).

To answer the question of what impacts AI will have on marketing practice, the authors applied a “marketing mix” approach, evaluating how AI can influence factors that play into consumer decision-making. These general categories are summarized in **Table A.12**.

Product	Price	Promotion (Brand)	Place (Sales & distribution)
<ul style="list-style-type: none"> • New product development • Hyper-personalisation • Automatic recommendations • Creating additional value • Additional solutions beyond product category 	<ul style="list-style-type: none"> • Price management and dynamic price matching to customer profile 	<ul style="list-style-type: none"> • Creating a unique experience • Personalised communication • Creating the wow factor and offering benefits • Elimination of the process of learning product categories • Positive impact on the customer • Minimised disappointment 	<ul style="list-style-type: none"> • Convenient shopping • The faster and simpler sales process • 24/7 customer service (chatbot) • Purchase automation • Service-free shops • Consultant-less customer support • New distribution channels • Merchandising automation

Table A.12. Areas of the Impact of AI on Marketing Mix

The authors conclude by stating that AI is certainly applied in many marketing areas, and the adoption of the AI techniques explored varies significantly. For example, while the modeling of decision-making patterns is widely adopted, voice recognition implementations are generally currently limited to products from companies like Google and Amazon. The authors also postulate that there is uncertainty on the part of some market practitioners about the true value of AI-enabled marketing, which perhaps contributes to the relatively slow pace of adoption in some areas. They also argue in passing that AI techniques are leading to better consumer experiences such as 24/7 support and more convenient shopping. Lastly, the authors argue that their work shows that often the adoption of AI within the marketing practitioner world is largely driven by the initial experimental adoption of one small solution rather than full adoption, and they propose that this indicates the need for further research in this entire area. This work can read as somewhat conclusory, but the collected examples of AI applications, both theoretical and real-world, are instructive.

How Artificial Intelligence Will Change the Future of Marketing^[17]

To begin this article, the authors propose theoretical examples of how AI will change business models, sales processes, and customer services methods.

Firstly, they propose that the adoption of AI-driven or AI-enabled vehicles will fundamentally alter the transportation industry by impacting ride-sharing businesses, auto insurance, reducing drunk driving incidents, and shortening commute times. The authors even suggest that the adoption of AI-enabled vehicles will impact the value of real estate as commuting could become a less painful experience.

Secondly, the authors propose that the sales industry will be impacted through the eventual adoption of AI-based natural language processing systems which operate in real time to monitor conversations between sales staff and potential customers, offering advice on how to guide the conversation to a sale. The authors also claim that eventually businesses may

wish to use AI-based agents to make initial customer contacts but acknowledge that this runs the risk of making customers uncomfortable if they discover they are not talking to a real person.

Thirdly, and perhaps most convincingly, the authors explore the application of AI to businesses operating on a “shipping-then-shopping” model. In this approach, the authors argue that AI can be employed to learn and predict what specific consumers want and are most likely to pay for, shipping these items without a formal order and subsequently accepting payment for the items kept by the consumer.

The authors also provide a table of current firms working in these three spaces as well as specific AI applications (**Table A.13**).

Industry or Usage Context (specific firm or AI application)	Description
AI in driverless cars (e.g., Tesla)	In the future, AI-enabled cars may allow for car journeys without any driver input, with the potential to significantly impact various industries (e.g., insurance, taxi services) and customer behaviors (e.g., whether they still buy cars).
Online retailing AI (e.g., Birchbox)	AI will enable better predictions for what customers want, which may cause firms to move away from a shopping-then-shipping business model and toward a shipping-then-shopping business model.
Fashion-related AI (e.g., Stitch Fix)	AI applications support stylists, who curate a set of clothing items for customers. Stitch Fix’s AI analyzes both numeric and image/other non-numeric data.
Sales AI (e.g., Conversica)	AI bots can automate parts of the sales process, augmenting the capabilities of existing sales teams. There may be backlash if customers know (upfront) that they are chatting with an AI bot (even if the AI bot is otherwise capable)
Customer service robots (e.g., Rock’em and Sock’em; Pepper)	Robots with task-automating AI respond to relatively simple customer service requests (e.g., making cocktails).
Emotional support AI (e.g., Replika)	AI aims to provide emotional support to customers by asking meaningful questions, offering social support, and adjusting to users’ linguistic syntax.
In-car AI (e.g., Affectiva)	In-car AI that analyzes driver data (e.g., facial expression) to evaluate drivers’ emotional and cognitive states.
Customer screening AI (e.g., Kanetix)	AI used to identify customers who should be provided incentives to buy insurance (and avoid those who (1) are already likely to buy and (2) those unlikely to buy).
Business process AI (e.g., IBM Interact)	AI used for multiple (simple) applications, such as customized offers (e.g., Bank of Montreal).
Retail store AI (e.g., Café X, LoweBot, 84.51, Bossa Nova)	Robots that can serve as coffee baristas, respond to simple customer service requests in Lowe’s stores, and identifying misshelved items in grocery stores.
Security AI (e.g., Knightscope’s K5)	Security robots patrol in offices or malls, equipped with superior sensing capabilities (e.g., thermal cameras).
Spiritual support AI (e.g., BlessU-2; Xian’er)	Customizable robot priest/monk offering blessings in different languages to the user.
Companion robot AI (e.g., Harmony from Realbotix)	Customizable robot companion, which promises reduced loneliness to the user.

Table A.13. Select Use Cases of AI

The authors next provide a high-level overview of artificial intelligence, touching on topics such as context awareness and levels of customer willingness to adopt the technology. The authors discuss prior research in the areas of task automation, robotics, and AI suitability

for different task types before moving into a discussion of the current state of AI and possible future evolutions. The authors provide several use cases for AI in the short and medium term, including price optimization, virtual assistants, and robots in retail (serving as baristas, for example).

The authors make suggestions for future research, including further work on AI and marketing strategy, sales, and customer adoption (engendering positive consumer views on AI). Lastly, the authors discuss potential policy issues, such as data privacy, bias within AI systems, and general ethical considerations. The authors conclude with a statement that while AI has already had a significant impact on marketing and business, this impact will only grow in the future. Overall, this work oscillates between well-considered passages and more speculative suggestions, but it retains value as an interesting exercise in current and theoretical AI applications in the business and marketing worlds.

Churn Prediction

Feature-selection-based Dynamic Transfer Ensemble Model for Customer Churn Prediction [18]

This work addresses accurate prediction of customer churn, examining related approaches and highlighting some of the core issues in this area of research. One of the central problems is that churn customers make up a very small percentage of the available data, while simultaneously being arguably the most valuable class to study within that data. This is referred to as “class imbalance,” a problem with the distribution of the data being analyzed. When this data is imbalanced, it makes the systems for predicting churn more likely to misclassify churn customers. The key approaches that experience this issue are decision trees, artificial neural networks, logistic regression, Bayesian classifiers, and support vector machines.

The authors propose an approach based on a combination of transfer learning, multi-classifier ensemble learning, and a “group method of data handling”-based neural network. The authors suggest that what they call a “feature selection-based dynamic transfer ensemble model” (FSDTE) should be applied to customer churn prediction. The proposed solution takes in data in the given domain as well as data from related domains to generate predictive models. The mathematical representation of the proposed approach is given, and empirical analysis of the prediction performance is conducted on two customer churn datasets. One dataset comes from a machine learning database at University of California Irvine based on cellular service provider customers, and the other dataset is taken from a credit card business and commercial bank in Chongqing, China.

The results of applying the proposed FSDTE model to these two data sets are evaluated, and the authors rigorously analyze the model, including variations in parameters, different training models, and changes in feature selection and classifiers. The authors conclude by stating that their proposed model deals well with class imbalance, and outperforms two traditional models for predicting churn, as well as three specified transfer learning models. The authors recommend further research to reduce the time complexity of their FSDTE

model and to address the fact that this model requires the same number of features in the source domain and target domain.

Negative Correlation Learning for Customer Churn Prediction: A Comparison Study^[19]

In this study, the authors propose the use of a multi-layer perceptron (a kind of neural network) trained on data obtained through negative correlation learning to predict customer churn within a telecommunications company. As seen in the previous paper, the authors here also discuss the problem of class imbalance in customer data sets, suggesting that the negative correlation learning approach can successfully tackle this challenge.

A brief discussion of related work follows, with the authors highlighting six previous works on customer churn prediction which use various machine learning and algorithmic models. A discussion of the mathematical formulations of a multi-layer perceptron and negative correlation learning follows, leading to the authors discussing the data set on which their validation experiments were performed. The dataset used was provided by a major Jordanian cellular telecommunication company and consisted of 12 attributes or features of 5000 randomly selected customers collected for a three-month period. The attributes are summarized in **Table A.14**.

Attribute name	Description
3G	Subscriber is provided with 3G service (yes, no)
Total consumption	Total monthly fees (calling + sms) (JD)
Calling fees	Total monthly calling fees (JD)
Local sms fees	Monthly local sms fees (JD)
International sms fees	Monthly fees for international sms (JD)
International calling fees	Monthly fees for international calling (JD)
Local sms count	Number of monthly local sms
International sms count	Number of monthly international sms
International MOU	Total of international outgoing calls in minutes
Total MOU	Total minutes of use for all outgoing calls
On net MOU	Minutes of use for on-net-outgoing calls
Churn	Churning customer status (yes, no)

Figure A.14. Attributes in the Telecommunications Data Set

The experimental results were evaluated on the criteria of the rate of correct classification, the rate of predicted churn, and the actual churn rate. The results showed that the proposed scheme's accuracy ranked among the top five techniques of the fifteen it was compared with and it is similarly ranked on its predicted churn rate and the actual churn rate. The authors conclude that the application of a multilayer perceptron trained with negative correlation learning data is a viable approach to predicting customer churn, as it outperforms some competitive machine learning approaches in the relevant literature.

Customer Interaction

Ascend by Evolv: Artificial Intelligence- Based Massively Multivariate Conversion Rate Optimization^[20]

In this article, the authors discuss a commercially available solution called Evolv, which is designed to use evolutionary algorithms to enhance web interfaces with the goal of increasing the website's "conversion rate optimization," increasing the number of site visitors who take a desired action.

In discussing the design of web interfaces, the authors note that it is often challenging to determine how the arrangement of elements may impact user actions. For example, a headline placed in a particular position on one page may lead to successful user action but may have the opposite effect when placed in the same place on a different page. Given this, it is highly desirable for site-owners to position elements for optimal conversion. The interaction between elements is therefore very important to optimizing conversion rates, which is challenging for humans to analyze and model. A standard method of addressing this difficulty is called "A/B" testing in which users are shown two versions of a page and the conversion rates of each page are recorded and compared. However, even if it is possible to identify an element or page that is not performing as desired, it can be difficult to determine what alternative configuration would lead to better result.

Evolv approaches this problem through the application of evolutionary algorithms. Each webpage is conceived of as a "genome" which is then subjected to recombination of the elements with random "mutations" performed on some elements to generate new page candidates. After evaluation, the best performing candidate is selected.

The authors compare the evolutionary algorithm approach to another automated method, the Taguchi method of multivariate analysis. The authors create a simulation environment which included simulated traffic and conversion rates in which to test both approaches. For the evolutionary algorithm test, the outcome of each simulated user interaction is considered as an "evaluation" of the given design, and then an average is taken across the users to estimate the quality of the design. The Taguchi method tests variables in small numbers of combinations, and evaluation is performed based on the combinations that score the highest within the scheme's valuation metrics. The highest scoring variables are then selected to construct the "best" combination.

The authors performed three different tests under varying conditions, finding that the evolutionary approach outperformed the Taguchi method considerably, in part because it makes continual improvements to the design throughout the course of the experiment, while the Taguchi method creates a single set of candidates for testing.

Some interesting challenges arise when the evolutionary algorithm approach is implemented in the real world. Since this evolution is designed to take place in a live environment, it is necessary to show users both the good and bad designs, but without intervention the website traffic will be evenly split between good and bad designs, reducing the overall conversion rate. This problem is addressed by using a multi-armed bandit approach to intelligently allocate traffic. The multi-armed bandit approach was tested in a simulated environment and was demonstrated as a suitable solution to creating and maintaining a strong overall conversion rate even during evolution.

The authors conduct a brief overview of potential future work, including the suggestion of an “always-on” approach under which the site in question would constantly evolve, rather than being merely an implementation of what the “best” candidate was after a given period of evolution. In conclusion, the authors note that the Evolv scheme is the first automated system for this type of conversion rate optimization and suggest that in the future a similar framework could be extended to new areas.

Forecasting Artificial Intelligence on Online Customer Assistance: Evidence from Chatbot Patent Analysis^[21]

This work undertakes an analysis of chatbot patents from 2000 to early 2020. The stated aim of this research is to gain a holistic picture of the trends and progress made in the use of artificial intelligence in chatbots. The authors identify two key areas of patent advancement that appear in their findings: the improvement of a chatbot’s ability to draw inferences on users from multiple sources, and the use of consumer knowledge to allow the chatbot to provide more deeply customized solutions. The patent analysis is conducted with the aim of evaluating trends, with a particular focus on chatbots designed to interact with and provide assistance to consumers. The findings indicate that there has been a strong trend towards patents involving so-called “conversational agents” which are based on the use of natural language.

Prior research has identified that the more chatbots display “human” behaviors and attributes, the more comfortable customers are interacting with them. However, the authors note that insufficient research has been done on how improvements can be made to increase consumer acceptance and trust of digital agents of this kind. The authors characterize more “traditional” chatbots as having scant ability to make sense out of natural language, making them ill-suited to sensing the reactions of customers and the state of the relationship between the bot and the customer. However, this implementation of chatbots is still strong in applications involving simpler question and answer interactions. The authors identify the central challenge faced by those making advancements in this area as “how to effectively combine the input in individuals’ natural language and the output in the machine’s language.”

The bulk of this paper’s analysis comes from the set of 668 patents collected by the researchers from across classification domains, which were then evaluated through text-based analysis of the patent abstracts. This approach is described as using “text mining to deeply identify the thematic patterns within each patent document.” The authors’ findings showed a strong patent focus on conversation and conversational capabilities, as well as significant interest in natural language simulation to improve agent-consumer interactions. The main phrases identified in the relevant patents are given in **Table A.15**.

Identification of the main phrases.					
	FREQUENCY	NO. CASES	% CASES	LENGTH	TF • IDF
CONVERSATIONAL AGENT	127	19	2,84%	2	196,3
NATURAL LANGUAGE	89	37	5,54%	2	111,8
DIALOG SYSTEM	70	10	1,50%	2	127,7
AGENT SERVICE	50	17	2,54%	2	79,7
CONVERSATIONAL AGENT SERVICE	47	16	2,40%	3	76,2
COMPUTING DEVICE	45	19	2,84%	2	69,6
INVENTION RELATES	45	43	6,44%	2	53,6
SERVICE SERVER	45	18	2,69%	2	70,6
USER INTERFACE	43	23	3,44%	2	62,9
AGENT SERVICE SERVER	38	15	2,25%	3	62,7

Table A.15. Identification of the Main Phrases in Relevant Chatbot Patents

The authors conclude with the statement that their research uncovers limited but noteworthy progress toward the goal of conversationally capable chatbots. They further note that additional research is needed to determine the extent to which the advancements proposed by the relevant patents will actually impact chatbot-customer interaction.

Frontiers: Machines vs. Humans: The Impact of Artificial Intelligence Chatbot Disclosure on Customer Purchases ^[22]

In this work, the authors undertake an examination of what happens in a consumer-chatbot interaction when the consumer is first informed that the digital agent is not a real person. This study involved a field experiment conducted on more than 6,200 customers who were randomly chosen to receive a sales call from either a chatbot or a real human.

The authors identify that one of the main obstacles to the adoption of chatbots is customer discomfort. Consumers are hesitant to share personal information with digital agents and display a dislike for allowing chatbots to help guide purchase decisions. According to research cited by the authors, customers perceive bots as less trustworthy than their human counterparts, which makes it challenging for those who are using chatbots to decide whether to disclose that fact. The authors contrast the business-focused goal of maximizing the value of chatbots with the ethics-focused argument that consumers have a right to know whether

they are interacting with a human or a bot. The authors note that regulators are becoming progressively more concerned about privacy protection with regard to human-bot interactions, encouraging anyone who uses chatbots to reveal that fact to their customers.

The study undertaken here categorized the entities making the sales calls into four categories: underdog workers (new, inexperienced workers), proficient workers, AI bots without disclosure, and AI bots with disclosure. Upon analysis of the results after contacting more than 6,200 customers, the authors find that AI bots who disclose their status as bots have lower purchase rates, higher hang-up rates, and shorter call length. Those bots which did disclose their “identity” saw a 79.7% reduction in purchase rates in comparison to the AI chatbots which did not disclose.

The authors briefly describe some potential ways to reduce the negative effects of chatbot disclosure, including simply disclosing the identity of the chatbot at the conclusion of the conversation. Notionally, this might reduce negative feelings towards chatbots as the customer will theoretically have a positive interaction which is then revealed to have been with a bot, hopefully reducing distrust on the part of the customer for future interactions. The authors conclude with a statement that chatbot disclosure has a significant impact on purchase rates, and that future research in this area is strongly encouraged as the adoption of AI in this space continues to accelerate.

Business Applications of Neuro-fuzzy Systems

A Review on the Applications of Neuro-fuzzy Systems in Business^[23]

This extensive review examines the use of neuro-fuzzy systems (NFS) in literature between 2005 and 2015, identifying various business applications and research trends. The authors first address the use of NFS in finance, describing its use in applications such as trading, credit risk assessment, evaluating the impact of bad loans, and business failure prediction. The examined literature shows that there is notable sophistication in this area.

Regarding marketing and distribution, the literature shows varied applications of NFS ranging from price prediction to assessing the viability of a new product. One cited study used NFS to assess the flexibility of a given supply chain, and a number of other studies address problems of consumption estimation, as in the cases of natural gas and electricity. The authors next turn to literature on AI applications in human resources management. While they note that AI is used in some adjacent applications such as addressing staffing and training, only a small number of papers on direct HR applications have been published since 2005.

NFS has seen much wider adoption in the realm of production and operations, and the authors note that reliance on such approaches only seems to be gaining popularity. The authors point out that this area has received huge attention in the literature, with diverse applications such as robot control, supplier selection, and even the monitoring of tool wear being shown. One particularly novel application of NFS was the prediction of fabric wrinkle recovery in textiles.

NFS applications in the world of business planning are also discussed, but like the field of human resource management discussed above, the authors note a relative dearth of literature in this area. The authors do identify some applications in this area, such as using AI to address stressful real-time decision problems, but the literature in this area is scarce at best.

Lastly, the authors discuss the applications of NFS in the realm of information systems, noting uses such as evaluating the quality of software, estimating the cost of software, managing customer relationships, as well as other applications in the software development process. The authors discuss the results of this review, highlighting that neuro-fuzzy systems are indeed used in diverse business domains, for diverse purposes. The authors provide some easily digestible graphics describing the topic breakdown of the uncovered literature (**Figure A.5** and **Figure A.6**).




Figure A.5. Relative Impact of Neur-fuzzy Systems in Different Business Areas




Figure A.6. Yearly Distribution of Research Articles on Neuro-fuzzy System Business Opportunities

In conclusion, the authors state that research trends on the business applications of NFS across the years surveyed show a strong focus on continuously improving the performance of these systems, as well as a desire to apply these approaches to new business problems. The authors note that in a large number of the studies surveyed, NFS solutions performed better than many other AI approaches, especially with regard to convergence and prediction accuracy. The authors advocate for further work to be done in underrepresented areas such as HR management, accounting, and business planning, and state that neuro-fuzzy systems have proved and will continue to prove indispensable for myriad business applications.

Recommendations for Future Work

There is a great deal of future work that could be done to expand upon the research provided in this review. At the inception of this project, a large keyword set was generated to aid in the search for relevant material. The goal was to identify more specific keywords which would narrow search results down to those specifically applicable within the identified AI application. This approach unfortunately needed to be set aside due to time constraints, but future work could easily use this keyword set to guide another round of literature review, potentially identifying further areas worthy of the TAC's attention.

Conclusion

This project highlights the variety of AI approaches being applied to problems within the telecommunications space. There is rarely (if ever) an absolutely dominant approach for any given problem set. Solutions and alternatives are proposed, discussed, discarded, and adopted based on a wide spectrum of evaluation criteria, and the various proposed applications of AI discussed throughout the preceding sections exhibit differing levels of maturity. These applications will continue to mature, disappear, and evolve, and new

applications and techniques will continue to appear at a high rate of speed. Therefore, it would be of great benefit to the TAC's understanding of the state of the art to review the new, relevant academic material on a regular basis.

Acknowledgements

My thanks go out to TAC members Adam Drobot and Dale Hatfield, who both provided indispensable guidance with this project. Their assistance and expertise made this work possible. Additional thanks are due to Jane Thompson, Associate Director of Faculty Services and Research at the William A. Wise Law Library at the University of Colorado Law School for her clear and experienced advice about how to approach this project.

Bibliography

1. Machine Learning Paradigms for Next-Generation Wireless Networks
C. Jiang, H. Zhang, Y. Ren, Z. Han, K. Chen and L. Hanzo, "Machine Learning Paradigms for Next-Generation Wireless Networks," in *IEEE Wireless Communications*, vol. 24, no. 2, pp. 98-105, April 2017, doi: 10.1109/MWC.2016.1500356WC.
2. Artificial Intelligence for 5G and Beyond 5G: Implementations, Algorithms, and Optimizations
C. Zhang, Y. Ueng, C. Studer and A. Burg, "Artificial Intelligence for 5G and Beyond 5G: Implementations, Algorithms, and Optimizations," in *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 10, no. 2, pp. 149-163, June 2020, doi: 10.1109/JETCAS.2020.3000103.
3. Path Loss Prediction Based on Machine Learning: Principle, Method, and Data Expansion
Y. Zhang, J. Wen, G. Yang, Z. He, J. Wang, "Path Loss Prediction Based on Machine Learning: Principle, Method, and Data Expansion," in *Applied Sciences*, vol. 9, iss. 9, Jan. 2019, doi: 10.3390/app9091908
4. Facilitating Mesh Reliability in PPDR Scenarios Utilizing Artificial Intelligence
R. Pirmagomedov, D. Moltchanov, A. Ometov, K. Muhammad, S. Andreev and Y. Koucheryavy, "Facilitating mmWave Mesh Reliability in PPDR Scenarios Utilizing Artificial Intelligence," in *IEEE Access*, vol. 7, pp. 180700-180712, 2019, doi: 10.1109/ACCESS.2019.2958426.
5. A Survey of Machine Learning Techniques Applied to Self-Organizing Cellular Networks
P. V. Klaine, M. A. Imran, O. Onireti and R. D. Souza, "A Survey of Machine Learning Techniques Applied to Self-Organizing Cellular Networks," in *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2392-2431, Fourth quarter 2017, doi: 10.1109/COMST.2017.2727878.
6. Optimal and Scalable Caching for 5G Using Reinforcement Learning of Space-Time Popularities
Sadeghi, F. Sheikholeslami and G. B. Giannakis, "Optimal and Scalable Caching for 5G Using Reinforcement Learning of Space-Time Popularities," in *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 1, pp. 180-190, Feb. 2018, doi: 10.1109/JSTSP.2017.2787979.
7. Machine Learning Aided Context-Aware Self-Healing Management for Ultra Dense Networks with QoS Provisions
M. Qin, Q. Yang, N. Cheng, H. Zhou, R. R. Rao and X. Shen, "Machine Learning Aided Context-Aware Self-Healing Management for Ultra Dense Networks With QoS

Provisions," in *IEEE Transactions on Vehicular Technology*, vol. 67, no. 12, pp. 12339-12351, Dec. 2018, doi: 10.1109/TVT.2018.2877910.

8. From 4G to 5G: Self-organized Network Management Meets Machine Learning

J. Moysen, L. Giupponi, "From 4G to 5G: Self-organized Network Management Meets Machine Learning," in *Computer Communications*, vol. 129, pp. 248-268, Sep. 2018, doi: 10.1016/j.comcom.2018.07.015.

9. Multi-layered Intrusion Detection and Prevention in the SDN/NFV Enabled Cloud of 5G Networks using AI-based Defense Mechanisms

I. H. Abdulqadder, S. Zhou, D. Sour, I. T. Aziz, S. M. A. Akber, "Multi-layered Intrusion Detection and Prevention in the SDN/NFV Enabled Cloud of 5G Networks using AI-based Defense Mechanisms," in *Computer Networks*, vol. 179, Oct. 2020, doi: 10.1016/j.comnet.2020.107364.

10. Realtime Scheduling and Power Allocation using Deep Neural Networks

S. Xu, P. Liu, R. Wang and S. S. Panwar, "Realtime Scheduling and Power Allocation Using Deep Neural Networks," 2019 IEEE Wireless Communications and Networking Conference (WCNC), Marrakesh, Morocco, 2019, pp. 1-5, doi: 10.1109/WCNC.2019.8886140.

11. Distributed Learning Algorithms for Spectrum Sharing in Spatial Random Access Wireless Networks

K. Cohen, A. Nedić and R. Srikant, "Distributed Learning Algorithms for Spectrum Sharing in Spatial Random Access Wireless Networks," in *IEEE Transactions on Automatic Control*, vol. 62, no. 6, pp. 2854-2869, June 2017, doi: 10.1109/TAC.2016.2626578.

12. Artificial Intelligence-based Cognitive Routing for Cognitive Radio Networks

J. Qadir, "Artificial Intelligence Based Cognitive Routing for Cognitive Radio Networks," in *Artificial Intelligence Review*, vol. 45, pp. 25–96, 2016, doi: 10.1007/s10462-015-9438-6.

13. Primary User Characterization for Cognitive Radio Wireless Networks using Neural System Based on Deep Learning

D. López, E. Rivas, O. Gualdrón, "Primary User Characterization for Cognitive Radio Wireless Networks using a Neural System Based on Deep Learning," in *Artificial Intelligence Review*, vol. 52, Iss. 1, pp. 169-195, June 2019, doi: 10.1007/s10462-017-9600-4.

14. An Online Context-Aware Machine Learning Algorithm for 5G mmWave Vehicular Communications

G. H. Sim, S. Klos, A. Asadi, A. Klein and M. Hollick, "An Online Context-Aware Machine Learning Algorithm for 5G mmWave Vehicular Communications," in *IEEE/ACM*

Transactions on Networking, vol. 26, no. 6, pp. 2487-2500, Dec. 2018, doi: 10.1109/TNET.2018.2869244.

15. The Impact of AI on the Advertising Process: The Chinese Experience

X. Qin, Z. Jiang, "The Impact of AI on the Advertising Process: The Chinese Experience," in *Journal of Advertising*, vol. 48, iss. 4, pp. 338-346, Sep. 2019, doi: 10.1080/00913367.2019.1652122.

16. Marketing and Artificial Intelligence

K. Jarek, G. Mazurek, "Marketing and Artificial Intelligence," in *Central European Business Review*, vol. 8, iss. 2, pp. 46-55, 2019, doi: 10.18267/j.cebr.213.

17. How Artificial Intelligence Will Change the Future of Marketing

T. Davenport, G. Abhijit, G. Dhruv, B. Timna, "How Artificial Intelligence Will Change the Future of Marketing," in *Journal of the Academy of Marketing Science*, vol. 48, iss. 1, pp. 24-42, Jan 2020, doi: 10.1007/s11747-019-00696-0.

18. Feature-selection-based Dynamic Transfer Ensemble Model for Customer Churn Prediction

J. Xiao, Y. Xiao, A. Huang, D. Liu, S. Wang, "Feature-selection-based Dynamic Transfer Ensemble Model for Customer Churn Prediction," in *Knowledge and Information Systems*, vol. 43, iss. 1, pp. 29-51, Apr 2015, doi: 10.1007/s10115-013-0722-y.

19. Negative Correlation Learning for Customer Churn Prediction: A Comparison Study

A. Rodan, A. Fayyumi, H. Faris, J. Alsakran, O. Al-Kadi, "Negative Correlation Learning for Customer Churn Prediction: A Comparison Study," in *The Scientific World Journal*, vol. 2015, Mar 2015, doi: 10.1155/2015/473283.

20. Ascend by Evolv: Artificial Intelligence- Based Massively Multivariate Conversion Rate Optimization

R. Miikkulainen, M. Brundage, J. Epstein, T. Foster, B. Hodjat, N. Iscoe, et al., "Ascend by Evolv: Artificial Intelligence-based Massively Multivariate Conversion Rate Optimization," in *AI Magazine*, vol. 41, iss. 1, pp. 44-60, Spring 2020, retrieved from <https://search.proquest.com/docview/2411777042>.

21. Forecasting Artificial Intelligence on Online Customer Assistance: Evidence from Chatbot Patent Analysis

E. Pantano, G. Pizzi, "Forecasting Artificial Intelligence on Online Customer Assistance: Evidence from Chatbot Patent Analysis," in *Journal of Retailing and Consumer Services*, vol. 55, July 2020, doi: 10.1016/j.jretconser.2020.102096.

22. Frontiers: Machines vs. Humans: The Impact of Artificial Intelligence Chatbot Disclosure on Customer Purchases

X. Luo, S. Tong, Z. Fang, Z. Q, “Frontiers: Machines vs. Humans: The Impact of Artificial Intelligence Chatbot Disclosure on Customer Purchases,” in *Marketing Science*, vol. 38, no. 6, Sep 2019, doi: 10.1287/mksc.2019.1192.

23. A Review on the Applications of Neuro-fuzzy Systems in Business

S. Rajab, V. Sharma, “A Review on the Applications of Neuro-fuzzy Systems in Business,” in *Artificial Intelligence Review*, vol. 49, pp. 481-510, doi: 10.1007/s10462-016-9536-0.

Glossary^[1]

Artificial Neural Networks

Computing systems vaguely inspired by the biological neural networks that constitute animal brains. An ANN is based on a collection of connected units or nodes called artificial neurons, which loosely model the neurons in a biological brain. Each connection, like the synapses in a biological brain, can transmit a signal to other neurons. An artificial neuron receives a signal then processes it and can signal neurons connected to it.

Bayesian Inference

A method of statistical inference in which Bayes' theorem is used to update the probability for a hypothesis as more evidence or information becomes available.

Convolutional Neural Network

A convolutional neural network (CNN, or ConvNet) is a class of deep neural networks, most commonly applied to analyzing visual imagery.

Decision Trees

A decision tree is a decision support tool that uses a tree-like model of decisions and their possible consequences, including chance event outcomes, resource costs, and utility. It is one way to display an algorithm that only contains conditional control statements.

Deep Learning

Also known as deep structured learning, deep learning is part of a broader family of machine learning methods based on artificial neural networks with representation learning. Learning can be supervised, semi-supervised or unsupervised.

Deep Neural Network

A deep neural network (DNN) is an artificial neural network (ANN) with multiple layers between the input and output layers. The DNN finds the correct mathematical manipulation to turn the input into the output, whether it be a linear relationship or a non-linear relationship. The network moves through the layers calculating the probability of each output. Each mathematical manipulation is considered a layer, and complex DNN have many layers, hence the name "deep" networks.

Evolutionary Algorithm

A subset of evolutionary computation, an evolutionary algorithm (EA) is a generic population-based metaheuristic optimization algorithm. An EA uses mechanisms inspired by biological evolution, such as reproduction, mutation, recombination, and selection. Candidate solutions to the optimization problem play the role of individuals in a population, and the fitness function determines the quality of the solutions. Evolution of the population then takes place after the repeated application of the above operators.

Fuzzy Logic

A form of many-valued logic in which the truth values of variables may be any real number between 0 and 1 both inclusive. It is employed to handle the concept of partial truth, where the truth value may range between completely true and completely false.

Genetic Algorithm

A metaheuristic inspired by the process of natural selection that belongs to the larger class of evolutionary algorithms (EA). Genetic algorithms are commonly used to generate high-quality solutions to optimization and search problems by relying on biologically inspired operators such as mutation, crossover, and selection.

Learning Automata

A learning automaton is one type of machine learning algorithm studied since 1970s. Learning automata select their current action based on past experiences from the environment. It will fall into the range of reinforcement learning if the environment is stochastic and a Markov decision process (MDP) is used.

Long Short-Term Memory

Long short-term memory (LSTM) is an artificial recurrent neural network (RNN) architecture used in the field of deep learning. Unlike standard feedforward neural networks, LSTM has feedback connections. It can not only process single data points (such as images), but also entire sequences of data (such as speech or video). For example, LSTM is applicable to tasks such as unsegmented, connected handwriting recognition, speech recognition and anomaly detection in network traffic or IDSs (intrusion detection systems).

Markov decision process

In mathematics, a Markov decision process (MDP) is a discrete-time stochastic control process. It provides a mathematical framework for modeling decision making in situations where outcomes are partly random and partly under the control of a decision maker. MDPs are useful for studying optimization problems solved via dynamic programming and reinforcement learning.

Multi-armed Bandit

In probability theory, the multi-armed bandit problem (sometimes called the K- or N-armed bandit problem) is a problem in which a fixed limited set of resources must be allocated between competing (alternative) choices in a way that maximizes their expected gain, when each choice's properties are only partially known at the time of allocation and may become better understood as time passes or by allocating resources to the choice. This is a classic reinforcement learning problem that exemplifies the exploration–exploitation tradeoff dilemma. The name comes from imagining a gambler at a row of slot machines (sometimes known as "one-armed bandits"), who has to decide which machines to play, how many times to play each machine and in which order to play them, and whether to continue with the current machine or try a different machine.

Multilayer Perceptron

A multilayer perceptron (MLP) is a class of feedforward artificial neural network (ANN). The term MLP is used ambiguously, sometimes loosely to refer to any feedforward ANN, sometimes strictly to refer to networks composed of multiple layers of perceptrons (with threshold activation). Multilayer perceptrons are sometimes colloquially referred to as "vanilla" neural networks, especially when they have a single hidden layer.

Neuro-fuzzy Systems

In the field of artificial intelligence, neuro-fuzzy refers to combinations of artificial neural networks and fuzzy logic.

Online Learning Algorithms/Online Machine Learning

A method of machine learning in which data becomes available in a sequential order and is used to update the best predictor for future data at each step, as opposed to batch learning techniques which generate the best predictor by learning on the entire training data set at once. Online learning is a common technique used in areas of machine learning where it is computationally infeasible to train over the entire dataset, requiring the need of out-of-core algorithms. It is also used in situations where it is necessary for the algorithm to dynamically adapt to new patterns in the data, or when the data itself is generated as a function of time, e.g., stock price prediction.

Overfitting

The production of an analysis that corresponds too closely or exactly to a particular set of data and may therefore fail to fit additional data or predict future observations reliably. An overfitted model is a statistical model that contains more parameters than can be justified by the data. The essence of overfitting is to have unknowingly extracted some of the residual variation (i.e., the noise) as if that variation represented underlying model structure.

Q-learning

A model-free reinforcement learning algorithm to learn a policy telling an agent what action to take under what circumstances. It does not require a model (hence the connotation "model-free") of the environment, and it can handle problems with stochastic transitions and rewards, without requiring adaptations. For any finite Markov decision process (FMDP), Q-learning finds an optimal policy in the sense of maximizing the expected value of the total reward over any and all successive steps, starting from the current state. Q-learning can identify an optimal action-selection policy for any given FMDP, given infinite exploration time and a partly random policy.

Recurrent Neural Network

A recurrent neural network (RNN) is a class of artificial neural networks where connections between nodes form a directed graph along a temporal sequence. This allows it to exhibit temporal dynamic behavior. Derived from feedforward neural networks, RNNs can use their

internal state (memory) to process variable length sequences of inputs. This makes them applicable to tasks such as unsegmented, connected handwriting recognition or speech recognition.

Regression

A set of statistical processes for estimating the relationships between a dependent variable (often called the 'outcome variable') and one or more independent variables (often called 'predictors', 'covariates', or 'features').

Reinforcement Learning

An area of machine learning concerned with how software agents ought to take actions in an environment in order to maximize the notion of cumulative reward. Reinforcement learning is one of three basic machine learning paradigms, alongside supervised learning and unsupervised learning.

Reinforcement learning differs from supervised learning in not needing labelled input/output pairs to be presented, and in not needing sub-optimal actions to be explicitly corrected. Instead, the focus is on finding a balance between exploration (of uncharted territory) and exploitation (of current knowledge).

Supervised Learning

The machine learning task of learning a function that maps an input to an output based on example input-output pairs. It infers a function from labeled training data consisting of a set of training examples. In supervised learning, each example is a pair consisting of an input object (typically a vector) and a desired output value (also called the supervisory signal). A supervised learning algorithm analyzes the training data and produces an inferred function, which can be used for mapping new examples. An optimal scenario will allow for the algorithm to correctly determine the class labels for unseen instances. This requires the learning algorithm to generalize from the training data to unseen situations in a "reasonable" way.

Support Vector Machines

Supervised learning models with associated learning algorithms that analyze data used for classification and regression analysis. The Support Vector Machine (SVM) algorithm is a popular machine learning tool that offers solutions for both classification and regression problems.

Transfer Learning

A research problem in machine learning that focuses on storing knowledge gained while solving one problem and applying it to a different but related problem. For example, knowledge gained while learning to recognize cars could apply when trying to recognize trucks.

Unsupervised Learning

A type of machine learning that looks for previously undetected patterns in a data set with no pre-existing labels and with a minimum of human supervision. In contrast to supervised learning that usually makes use of human-labeled data, unsupervised learning, also known as self-organization, allows for modeling of probability densities over inputs. It forms one of the three main categories of machine learning, along with supervised and reinforcement learning.

[1] Definitions in this section are derived from Wikipedia.

8 Appendix B: AI/ML in Federally Funded Initiatives

The private sector is making substantial investments in the development and application of Artificial Intelligence (AI) in order to improve operational capability and to better serve their customers. The United States Government is doing its part to enhance the growth of the U.S economy and to maintain global leadership in AI by funding a wide variety of initiatives designed to advance research and development in AI. Appendix B summarizes some of the most important federally funded initiatives that have a bearing, either directly or indirectly, on the provision and demand for telecommunications and information services, as well as on the security of the systems that provide such services. [1] The selected initiatives fall within six distinct research and development areas.

8.1 Long Term and Targeted Investments in Research and Development

The National Institute of Standards and Technology (NIST) has initiated the Fundamental and Applied Research and Standards for AI Technologies (FARSAIT) program. This program examines methods for improving trust in AI systems by measuring and improving the security, explainability, and transparency of AI systems. NIST scientists are also examining the challenges associated with making wireless networks more adaptive and innovative. NIST believes that data science and artificial intelligence (AI) will be playing a more important role in the operation of many of the components of a wireless network. According to NIST the combination of data science and AI will together drive improved network performance, spectrum efficiency, and power consumption. NIST is examining how best to deal with the sometimes massive amount of data that needs to be transferred and crunched in order to fully leverage the network performance benefits of AI.

An important element in enforcing spectrum sharing agreements involves identifying, protecting, and classifying wireless signals. The Institute of Telecommunications Sciences (ITS) of the National Telecommunications and Information Administration (NTIA) is examining the use of Machine Learning (ML) and neural networks to identify, protect, and classify wireless signals, as well as to monitor spectral activity for abnormal or suspicious behavior. The Department of Homeland Security (DHS) is examining the use of AI technologies and applications for use in such areas as cybersecurity, counterterrorism, border security, disaster response, and the protection of critical infrastructure.

8.2 Human-AI Interactions and Collaboration

Artificial Intelligence employs, among other things, linear algebra to find the solution to a set of equations and calculus to optimize a chosen objective function (e.g., maximize prediction accuracy) given that solution. Both tools are employed to make a statement about something based on a set of data collected in the real world. However, given the complexity and sometimes non-intuitive nature of the relationships contained in the data, the actual reasoning behind the predictions rendered by an AI system may not be easily explained and

described. The absence of so-called “explainability” gives rise to concerns about the level of trust that should be assigned to a statement about that data. Concerns about trust have lead Department of Defense/DARPA to launch the Explainable Artificial Intelligence (XAI) program the objective of which is to find ways in which AI systems can better explain their reasoning in human understandable terms.

8.3 Ethical, Legal, Societal Implications of AI

The U.S. Government has initiated funding for programs that explore the ethical, legal, and societal implications of AI. Among the more important concerns is the issue of unrecognized adverse biases in the determinations made by AI systems. To that end, the National Science Foundation (NSF) has allocated funds to projects that focus on techniques that offer greater transparency and accountability, while at the same time mitigating adverse biases. At the same time, the Department of Homeland Security’s (DHS) S&T Data Analytics Technology Center (DA-TC) has focused its attention, using both qualitative and quantitative analyses, on bias and fairness in the results of AI systems. DA-TC is also funding research on the risk associated with the use of AI systems by malicious actors. In response to the increasing level of automation in America’s transportation system, the Department of Transportation’s (DOT) Federal Highway Administration’s Exploratory Advanced Research (EAR) program is exploring the development of artificial intelligence and machine learning technology within the surface transportation sector for the purpose of making transportation safer. For example, the EAR program is supporting AI research in the collection of traffic data to spot trends and identify relationships between seemingly disparate data streams.

8.4 Safety & Security of AI Systems

It is recognized that outside parties can strategically manipulate and mislead AI systems and, in so doing, cause substantial harm to telecommunications networks, autonomous vehicles, and financial systems, among other important systems. The Federal Government believes that it is imperative that every effort be made to discover and implement techniques and measures that will enhance the security and safety of AI systems.

Cybersecurity has become a major priority in nearly every Federal agency. In many instances, researchers are employing AI systems to help predict, identify, and deter cyberattacks. NIST researchers are investigating how AI systems can be used to automate vulnerability assessments for digital infrastructure and to produce vulnerability ratings using the industry-standard Common Vulnerability Scoring System. Complicating the task of pinpointing and fixing vulnerabilities of AI systems is the challenge posed by the non-transparent (e.g., black box) nature of such systems. To examine the issue of “explainability,” in January 2019 the National Science Foundation (NSF) and several leading universities established the “Center for Trustworthy Machine Learning,” a five-year effort to examine this and other security related issues. In addition, as noted earlier,

DARPA's XAI program is exploring ways to make AI systems more understandable for humans and, in so doing, reducing the vulnerability of AI systems to adversarial behavior.

In an attempt to inform future standards and best practices for assessing and managing vulnerabilities of ML technologies, the National Institute of Standards and Technology's (NIST) Information Technology and National Cybersecurity Center of Excellence has been developing models depicting the different processes parties employ to execute an adversarial ML attack. The Intelligence Advanced Research Project Activity's (IARPA) "Trojans in Artificial Intelligence" (TrojAI) program is funding projects to detect if an adversary has inserted a Trojan or backdoor into an AI system for purposes of, for example, "poisoning" the process by which an AI algorithm is trained.

Among the Federal agencies leading the way in conducting research in identifying cyberattacks are NSF and IARPA. In 2016, IARPA launched its Cyber-attack Automated Unconventional Sensor Environment program, which seeks to develop improved methods for both forecasting and detecting cyberattacks. NSF's Secure and Trustworthy Cyberspace program and IARPA's Virtuous User Environment program are examining ways in which AI can be used to improve detection of cyber anomalies. An important element of improving cyber security involves establishing an effective defense against cyberattacks. To that end, DARPA, IARPA, NIST, and NSF all have funded multiple research projects involving use of AI for cyber defense. For example, DARPA's Guaranteeing AI Robustness Against Deception program is researching how to build AI systems that are resistant to adversaries' efforts to fool them. NIST researchers are investigating how AI systems can be used to automate vulnerability assessments for digital infrastructure and to produce vulnerability ratings using the industry-standard Common Vulnerability Scoring System.

In some cases, a Federal agency acts as a coordinator for the development and implementation of AI tools. For example, the Department of Energy coordinates research projects that develop AI capabilities to enhance the cybersecurity of critical energy systems. Among other things, these research projects include a requirement to verify and validate emerging AI tools and technologies. For example, to lower the risk of inadvertent disruption in energy delivery, asset owners/operators provide access to testbeds that accurately replicate their operational infrastructure, allowing for validation of the developed technology without compromising safety and reliability.

8.5 Data Sets for Training & Testing AI Systems

Machine Learning is a process by which a computer discovers for itself, as opposed to being directed by specific logic written by a programmer, about how to make better decisions (e.g., predictions) over time based upon analysis it completes using data to which it has access. Under machine learning, additional insights obtained from new data leads to the computer's improved ability to make good decisions. Given the importance of data to the machine learning process and with it, the training and testing of AI systems, the curation of

and access to data represents a critical element of the Federal Government's efforts to maintain global leadership in AI.

Numerous Federal agencies have made important strides in creating data that are formatted in a manner that is amenable for both the training and testing of AI systems by private, academic, and government participants. These strides cover a wide variety of areas, including health, science and engineering, justice and security, and resource management. For example, NIST's Spectrum Sharing Models and Measurement Tools project is developing a curated radio frequency (RF) signal database to aid in the development of ML models for signal detection, identification, and classification. Example datasets include radar signals like those planned for the 3.5 GHz band, involving the level of noise and interference, data sets that can be used to train and evaluate AI-based detectors enabling Federal-commercial spectrum sharing.

NTIA is building a wireless testbed for monitoring and characterizing spectrum usage and is collaborating with other organizations to develop and implement a standard data specification for spectrum data. In 2019, the system consisted of six spectrum sensors distributed over a neighborhood-sized area that will continuously monitor a selected frequency band and store all collected data for use in research and analysis. Once operational, the spectrum monitoring database will be made publicly available for the purpose of enabling researchers to apply ML and other AI techniques to these extremely large datasets. For example, the database will provide researchers the opportunity to apply ML and other AI techniques to the problem of achieving dynamic coordination of shared spectrum access among Federal and non-Federal spectrum users.

NTIA also maintains collections of audio data gathered to evaluate human perceptions of audio quality where the human voice is masked by background noise, such as sirens, saws, and gunshots. These data, which are publicly available, provide researchers the opportunity to apply AI techniques to extract speech from audio streams. NTIA also maintains publicly available data involving video streams which were assessed by human subjects for their video quality. The data provide researchers the opportunity to train AI systems for the purposes of predicting the level of human perception in a given video environment.

The Federal Highway Administration (FHWA), a division of the Department of Transportation (DOT), collects a large amount of geometric, traffic and crash data. Housed in its Highway Safety Information System, the FHWA has begun research on using machine learning to find patterns and connections between different types of data (e.g., crash reports, hospital records). The DOT Secure Data Commons is a cloud-based analytics platform that enables traffic engineers, researchers, and data scientists to access transportation-related datasets. It provides a secure platform for sharing and collaborating on research, tools, algorithms, and analysis involving sensitive datasets using commercially available tools.

NIST has created shared datasets and environments spanning robotics, material discovery, and wireless spectrum analysis, among other areas. As noted earlier, NIST's Spectrum Sharing Models and Measurement Tools project will create a curated radio frequency (RF) signal database to aid in the development of machine learning models for signal detection and classification. In addition to developing new data sets, NIST is making investments in facilities that will speed up the machine learning training process and the testing of AI systems. For example, in 2018, NIST acquired a mid-scale computing facility optimized for big data and machine learning applications. One of the goals of this investment is to explore, in real time, the use of machine learning for the design and optimization of experimental systems.

8.6 Measuring & Evaluating AI System Performance

Measuring and evaluating AI system performance is a critical part of implementing AI and ML techniques. An important part of these processes is the establishment of metrics and benchmarks for evaluating the capabilities and risks of AI systems. NIST stands at the center of the Federal Government's efforts in measurement and benchmark work. In 2019 NIST released a report that identifies nine areas of AI standards: concepts and terminology, data and knowledge, human interactions, metrics, networking, performance testing and reporting methodology, safety, risk management, and trustworthiness. Further, NIST's Fundamental and Applied Research and Standards for AI Technology (FARSAT) program is supporting several research projects related to the development of standards and benchmarks, including projects to assess the performance of generative adversarial networks (GAN), an advanced form of Deep Learning, to measure the vulnerability of AI image-recognition tools to adversary attacks. The Department of Energy (DOE) researchers are employing ML to aggregate large operational datasets in order to benchmark performance and identify areas where current predictive models fall short.

End Note [1] The following borrows heavily from a document prepared by the National Science and Technology Council entitled "2016-2019 Progress Report: Advancing Artificial Intelligence R&D," November 2019. <https://www.nitrd.gov/pubs/AI-Research-and-Development-Progress-Report-2016-2019.pdf>.

9 Appendix C: DARPA's Spectrum Collaboration Challenge: A Review

9.1 Introduction

Military units must coordinate their activities in order to maximize both their defensive and offensive capabilities.[1] Such coordination extends critically to the use of spectrum given its importance for communication and the overall operational capabilities of a military unit. In 2016 the Defense Advanced Research Projects Agency (DARPA) initiated a program called the Spectrum Collaboration Challenge (SC2). The objective of the SC2 program was to motivate both private sector and academic research in finding new and improved ways in which users can share spectrum.[2] To that end, DARPA funded the development of a machine that attempts to emulate the naturally occurring world's electromagnetic environment. The purpose of the "emulator" was to serve as a controlled, experimental testbed in which different approaches to spectrum sharing can be evaluated. To motivate effort, the SC2 program included a contest in which teams competed over a series of years for \$17 million in prize money based on the team's spectrum sharing performance.

DARPA's own research program examined the spectrum sharing performance properties of a highly decentralized system in which an administrator plays a very limited spectrum management role.[3] Performance, broadly speaking, was measured by the extent to which a set of spectrum users (i.e., teams), in response to a set of rules established by the administrator, can over a period of time share spectrum so that teams were able to achieve their traffic delivery goals. Given the short time period over which sharing can occur, the research focused on whether the application of Artificial Intelligence (AI) and Machine Learning (ML) can guide the spectrum use demands of users in a manner that would allow teams to reach such goals.[4]

9.2 An Experimental Testbed – The Colosseum

Central to the research program is the development of an experimental testbed which can emulate tens of thousands of possible communications among a wide variety of wireless devices across a wide variety of spectral and operational environments. Built by National Instruments (NI) and assembled and maintained during the course of the SC2 by the John Hopkins Applied Physics Laboratory (APL), the facility (commonly referred to as the "Colosseum") includes 128, two-antenna, software defined radio (SDR) units. Each SDR contains both a graphical processing unit (GPU) and a central processing unit (CPU). APL worked with SI to include in the Colosseum 64 Field Programmable Gate Arrays (FPGA) to emulate the necessary set of electromagnetic waves associated with a given test environment.[5] As structured the Colosseum enabled the administrator to generate a variety of spectrum and operational environments defined by, among other things, the amount of spectrum available for use at any moment in time and space, the presence or non-presence of incumbent spectrum users, as well as the level of demand a given user would like to place

on the available spectrum.[6] A team's wireless network within the Colosseum consisted of a collection of radios designed and built by each team spread over a given geographic space. The Colosseum provided teams with a cloud-based platform with which to analyze data and employ machine learning to train their spectrum access algorithms.

9.3 The Spectrum Collaboration “Game”

The SC2 program involved the establishment of a set of spectrum scenarios defined by the amount of spectrum available for use, a set of five radio networks, one for each team, the successful operation of which required sufficient access to that spectrum, and a scoring system that assigns points to teams based on the degree of collaborative behavior which they exhibit during the course of given collaboration game.[7] Prior to every game, the administrator assigned a collection of “characteristics” to each team.[8] Among the characteristics are the GPS locations of the 10 radios that constitute the team's wireless network, a stream of data traffic (termed a “flow”) that the team wishes to send between specific radios within its own network, and a set of Quality of Service (QoS) requirements that are associated with each data flow.[9] Each spectrum environment consisted of 40 MHz or less of available spectrum, which teams can access over time based upon their own channelization needs.[10] Each team faced the challenge of employing machine learning (ML) techniques to develop a set of AI algorithms designed to predict on a real-time basis the demands that other users will place on spectrum.[11]

The administrator's scoring rules consisted of two main components. One component involved a performance threshold for each flow that the administrator assigned (i.e., “flow performance threshold”) to the team. If the flow didn't satisfy this performance requirement, the team would earn zero points for that flow. In addition to a performance threshold for each flow, there was also an aggregate scoring threshold that applied to the sum of the scores earned by the team for the respective flow performances.[12] A team's aggregate score had to exceed the aggregate scoring threshold in order for it to earn the scores it earned on each flow. The scoring rules placed an emphasis on both collaborative and competitive behavior. To induce spectrum sharing among teams the rules of the game were such that the “score” any one team received from a given level of spectrum use depended, at times, on the level of the resource that other teams use. In particular, the score achieved by a given team would be equal to the score earned by the lowest scoring team if the latter's score was below the aggregate threshold.[13] On the other hand, teams had an incentive to grab spectrum and, thus behave competitively, in instances in which all other teams had scores in excess of the threshold in an effort to satisfy the QoS requirements for a given flow assigned to them by the administrator.

To facilitate collaborative behavior among teams, teams were given the opportunity to use a dedicated channel to communicate to other teams their desired spectrum use and, in some cases, to communicate its estimate of its own score. Teams were required to report

information about their spectrum usage, network location, achieved performance level of all of their flows, and an estimate of their own aggregate score.[14] A procedure was established to promote the truthful reporting of such information.[15]

9.4 Impediments to Collaboration

There were several features of the spectrum access game that made collaboration difficult. First, the “game” in which teams participated was one of “incomplete information.” For example, the administrator did not reveal prior to the start of the game the level of the aggregate threshold used in the scoring process. Not knowing the level of the aggregate threshold at the start of the game was rendered somewhat moot, however, given that the administrator often changed the value of the threshold during the course of the game.[16] Second, because of the uncertainty related to the level of the threshold applied to each game, teams were uncertain about the level of their own scores during the course of the game. Third, despite the availability of the communication channel, teams did not know with certainty the spectrum access strategies of other teams, including the waveforms they wanted and the protocols they would employ during the course of the game. Fourth, teams were uncertain about the scores of other teams, including whether a given team’s score was below the threshold used for scoring. Fifth, teams were not given the opportunity to express their willingness to substitute one wavelength for another wavelength. The absence of such expressed flexibility, all things being equal, increased the likelihood that the spectrum demands of teams will overlap and, thus, reduced the level of achieved collaboration. Sixth, the large number of teams that participated in the contest combined with the shortness of the period during which a team’s algorithms could be trained made it difficult to generate enough data to effectively employ ML techniques for the purpose of developing an algorithm that performed well in predicting the spectrum access strategies of other teams.

9.5 Team Performance and Overall Results

The administrator calculated, during the course of the collaboration game, the degree of shared use of the electromagnetic waves that resulted from the demand each team placed on the available spectrum. The administrator also calculated scores for each team based on the level of their collaborative behavior, as well as their ability to achieve the quality of service requirements associated with their assigned flows. Teams received their actual scores at the end of each game.

Because of differences in team member composition and skill sets, teams varied substantially in the quality of the radios they built.[17] The SC2 lasted for approximately 3 years, during which the administrator altered the rules of the game in response to contest data that demonstrated, for example, an insufficient level of spectrum sharing among the teams. During this period of time, teams adjusted their spectrum access strategies, moving away from a pure “rules-based” strategy, to a strategy that is at least partially involved machine learning-based decision making.[18]

How successful was the SC2? One important objective of the challenge is to demonstrate that a highly decentralized spectrum management process is superior to a centralized spectrum management process as measured by the level of spectral efficiency achieved. DARPA reports that their proposed decentralized spectrum management system is superior to a centralized spectrum management process as measured by the amount of data that a group of five networks can transmit over a period of time. That is, five networks collectively transmitted more data under DARPA's decentralized spectrum management process than under a centralized process where an administrator assigned spectrum to those same networks on an exclusive use basis.[19] However, extending such results to the naturally occurring world must be done with great care. For example, a careful analysis requires modeling the process by which an administrator in the naturally occurring world would assign spectrum to the various teams on an exclusive basis, including how much spectrum is assigned to each user. Because it doesn't appear that such modeling was completed, substantial care must be used in extending the outcome of DARPA's comparative analysis to the naturally occurring process. This assessment points to the need for additional analyses regarding the comparative performance properties of a centralized versus a highly decentralized spectrum management process.

Another important objective involved demonstrating that ML-trained algorithms can contribute greatly to the superior spectral efficiency performance of a decentralized spectrum management process.[20] The results of the challenge demonstrate that the source of the improvement in spectral efficiency cannot be significantly attributed to the use of ML-trained algorithms. Indeed, the two best performing teams acknowledge that ML-trained algorithms were of limited use in assisting them in improving their individual scores. The reasons are several fold. First, as noted, the resource allocation problem, even viewed in static terms, was very complex. Second, the large strategy space available within a game, and the associated lack of data with which to train an algorithm, limited the usefulness of AI in enhancing collaboration among teams.

Given the state of the science regarding spectrum sharing, the SC2 program's stated objectives were very ambitious. And because of that, it is understandable that the program was not able to fully achieve its two important objectives. Much additional work on the science of spectrum sharing needs to be completed. That work can partially rest on several of the results obtained in the SC2 program. First, the results of SC2 program demonstrate that the teams were able to achieve, in a dynamic setting and across a wide variety of spectral and operational conditions a level of collaboration that resulted in spectrum sharing. Second, the results also demonstrate that the rules of a spectrum sharing game matter. In particular, because of the rules of the game, teams had an incentive at times to jam the transmissions of other teams in an effort to reduce the scores earned by those teams. Such non-cooperative behavior served to reduce the level of spectral efficiency achieved by the process. Third, the results also demonstrate that AI and machine learning, rather than

driving collaboration results, assisted teams in identifying private-welfare maximizing spectrum sharing strategies measured over a very short time scale (i.e., seconds). This result points to the need for additional work to further explore the role that AI and machine learning can play in enhancing spectral efficiency. Finally, the results hint at the important role that an RF emulator testbed may play in improving the spectrum management process generally.[21] An RF emulator includes a mathematical model that attempts to replicate the technical characteristics of a naturally occurring spectrum environment. A well-designed and highly scalable emulator may be used to test, in a controlled experimental setting, hypotheses about the effects on interested parties of proposed changes in the spectral environment.[22] The formal, objective testing of hypotheses (e.g., the effect of a change in a spectrum allocation on the performance of an existing wireless network) may, by virtue of the data generated, improve the efficiency and quality of the decisions made by policymakers.

9.6 Extending DARPA's Spectrum Sharing Work for Commercial Users

The SC2 program sheds light on the ability of a highly decentralized spectrum management process to induce a group of spectrum users to share spectrum designated for non-exclusive use. Nonetheless, the Commission will likely be hard pressed to adopt DARPA's approach for the sharing of spectrum in a commercial environment.[23] At a very basic level, the efficient use of a spectrum, as with other resources, requires sorting parties into two categories. In the case of spectrum, one category is the set of parties that have access to spectrum, while the other category includes the set of parties that do not have access to spectrum.[24] The SC2 program uses two tools to sort users into the two categories. One tool is a communication channel that enables users to communicate their intention to use one or more channels over a period of time in the hope of avoiding too much use on those channels. Another tool involves the use of a scoring system and the financial rewards that teams obtain within that system.

Each tool has its weaknesses. For example, teams can use the communication channel to mislead others about their true intentions in the hope that such misrepresentation would conceal the spectrum the team truly needs to satisfy its quality of service requirements. The second tool also has its weaknesses. The adopted scoring system induced cooperative as well as, in some situations, non-cooperative behavior among teams. It is well understood that "players" (e.g., Teams) typically have a difficult time finding a suitable equilibrium in such games when as few as two players are involved, let alone a game in which five players participate.[25] Importantly, the economy is forced to address similar resource allocation problems all the time. It attempts to solve the problem of sorting users and non-users out by the use of market prices. In particular, the "efficient" market price is one that correctly identifies those users that have the greatest need for the resource, based upon their expressed willingness to pay for that resource, and compares the value of those needs to the cost society incurs from satisfying such needs. Market prices provide this sorting function by

inducing market participants to reveal information regarding the extent to which they wish to have access to (or offer to provide) a resource. The efficient allocation of a given resource is achieved when there are no changes in the allocation of that resource across users that will improve the value that society places on that resource.[26]

An alternative approach to addressing the spectrum allocation problem involving shared spectrum use involves incorporating a price mechanism to differentiate spectrum users from non-users of spectrum in instances where there is excess demand.[27] This approach has been relied on in some cases (e.g., CBRS) where spectrum users have both the right to employ spectrum on an exclusive basis and, further, the right to sell rights to use their spectrum to others. But in such cases, market participants employ bilateral bargaining to complete a transaction. Moreover, in each of these cases, the time period for which a transaction applies is measured in days as opposed to seconds. Uncertainty exists among prospective participants about whether bilateral bargaining will lead to the efficient sharing of spectrum. Indeed, economic science has demonstrated that bilateral bargaining often results in economically inefficient outcomes which, in this case, means the inefficient sharing of spectrum.

9.7 Moving Beyond Bi-lateral Bargaining

Is it possible to jettison reliance on a bilateral bargaining process to accomplish spectrum sharing? Could that process accomplish spectrum sharing over intervals of time less than one month? Answering these important questions requires the careful application of engineering and economic science. The research begins by adopting certain elements of DARPA's approach in solving the resource allocation problem involving shared spectrum use. For example, DARPA proposes a "centralized" process in which users submit their spectrum use demands to a single location more or less at the same time. One important feature of such centralization is that it opens up the possibility of evaluating the needs of all users simultaneously. The simultaneous evaluation of user spectrum needs allows the process to not only identify the most important needs as measured by willingness to pay, but also assists in more completely satisfying those needs. Second, the DARPA's spectrum sharing process explores the usefulness of AI and ML in predicting the spectrum use demands of users where spectrum is made available over short intervals of time. The introduction of AI and ML may be a very important part in the development of a new and improved method for sharing spectrum. It also involves employing a testbed that would allow a clear analysis of the performance properties of the proposed mechanism across a wide variety of spectral and operational environments. The mechanism itself would include rules that identify acceptable market behavior by participants, a set of mathematical algorithms that can do the appropriate sorting of users given their expressed demands for spectrum, and the development of pricing algorithms that determines how much, if any, a spectrum user needs to pay in order to access a given band of spectrum over a period of time.[28]

[1] The following is based on publicly available material, as well as material derived from presentations made before the AI Working Group by teams representing the University of Florida (GatorWings) and Vanderbilt University (MarmotE).

[2] DARPA's vision involves the development of a collaborative intelligent radio network (CIRN) where radio networks will autonomously collaborate in order to better share the RF spectrum. For a discussion of DARPA's vision, see Paul Tilghman, "If DARPA Has Its Way, AI Will Rule the Wireless Spectrum," IEEE Spectrum, 28 May 2019. <https://spectrum.ieee.org/telecom/wireless/if-darpa-has-its-way-ai-will-rule-the-wireless-spectrum>.

[3] Contrary to standard spectrum management procedures, under the SC2 challenge the administrator neither established a predetermined number of licenses nor imposed a predetermined set of frequency and bandwidth allocations. Rather, under the SC2 challenge, frequency and bandwidth allocations (assignments) result from the spectrum access decisions made by the various teams. Further, unlike in CBRS where SAS (Spectrum Access System) plays an important and direct role in facilitating the sharing of spectrum, the SC2 challenge did not include any automated process for the sharing of spectrum.

[4] The time period over which decisions regarding spectrum use needed to be made, typically in seconds (versus in one day increments as in CBRS), required the use of autonomous agents.

[5] As noted by DARPA, "By the numbers, the Colosseum is a 256-by-256-channel RF channel emulator, which means that it can calculate and simulate in real-time more than 65,000 channel interactions among 256 wireless devices. Each simulated channel behaves as though it has a bandwidth (information content) of 100 MHz, which means the testbed supports 25.6 GHz of bandwidth in any instant. Moreover, each channel's transmission and reception frequency is tunable between 10 MHz (as in broadcast FM radio) and 6 GHz (as in WiFi)." See <https://archive.darpa.mil/sc2/news/worlds-most-powerful-emulator-of-radio-signal-traffic-opens-for-business>.

[6] The programmable nature of the Colosseum enabled the administrator to emulate a wide variety of spectral and operating environments (e.g., a densely populated city, an open field, a suburban shopping mall).

[7] The SC2 extended for approximately 3 years, during which time the administrator changed the rules of the collaboration game in an attempt to enhance the level of spectrum sharing. The rules presented in this section represent the final rules adopted by the administrator.

[8] The ability to assign different characteristics to teams created the opportunity to examine the degree to which sharing is sensitive to the level of diverseness among the teams with respect to the services they wish to provide. While each team's network consisted of 10 radio, the locations of the radios assigned by the administrator varied across teams. To simulate mobility, the administrator often changed the locations of team's radios during a given game.

[9] The administrator assigned to each team a data flow, and a quality of service (QoS) requirement for that data flow, for each segment of the team's network. The QoS requirements were defined in terms of throughput, latency, hold time, among other performance measures.

[10] The amount of spectrum available to share varied not only across games but also sometimes within a game.

[11] Among other things, the demand that any one team places on the available spectrum must also consider the need to protect incumbent users of spectrum and, in some cases, complications associated with jammers.

[12] The aggregate scoring threshold sometimes varied across teams within a given game.

[13] Given that a team's payoff may depend on the score achieved by another team, teams had a strong incentive to communicate to each other information about their respective channels needs and protocols.

[14] However, when a team's estimated aggregate score exceeded its estimate of the aggregate threshold, the team was not required to report to other teams its estimate of its aggregate score.

[15] For example, following the completion of a game the administrator checked the information provided by a team against the actual actions taken by the team. Such checking was part of the pre-qualification process established by the administrator for admittance into the final competition.

[16] During the course of the game the administrator provided each team information about the value of the performance threshold for each flow.

[17] The radios developed by the teams varied not only on such matters as tunability (e.g., noise rejection) and frequency agility, but also on basic radio design and architecture.

[18] A "rules-based" spectrum access strategy is one in which the user establishes a set of Boolean-based (e.g., if then statements) statements that remain fixed regardless of changing conditions in the spectrum use environment. According to reports, while the set of rules adopted by teams early in the challenge varied across teams, they shared some common features (e.g., if there are no empty frequency bands radios should select the ones with the least interference). Later in the challenge teams began to rely on AI to make predictions regarding what the status of the spectral environment will be in the future and, based upon those predictions, employed AI to develop a strategy to access the necessary spectrum on a going forward basis.

[19] The performance of the decentralized spectrum management process was negatively affected at times by the decision of teams to engage in "jamming," defined herein as a deliberate attempt by a team to lower another team's score by accessing spectrum needed by that team.

[20] DARPA suggested that ML-based algorithms could, by virtue of mathematical brute force, lead to substantial improvements in spectral efficiency.

[21] While wireless researchers are quite familiar with RF emulators to test the real-time performance of wireless devices and base stations, RF emulators are not typically employed by spectrum policymakers to resolve disputes between and among interested parties regarding the effects of possible changes in the spectrum environment.

[22] The ability of the emulator to carefully test a hypothesis depends on the quality of the mathematical model that emulates the desired naturally occurring environment. Disagreements may arise among interested parties regarding whether the model is sufficiently detailed so as to make reasonably good predictions about the naturally occurring environment. For example, there exists disagreement about whether the Commission should continue to employ the Longely-Rice model to predict transmission loss (i.e., frequencies between 20 MHz and 40 GHz and path lengths between 1 km and 2000 km) over irregular terrain.

[23] The reasons range from a likely hesitancy to reward commercial users of spectrum for cooperative behavior to the Commission's emphasis on relying on market processes (e.g., auctions) to determine the identity of parties that employ spectrum.

[24] As noted in earlier sections, and as shown by the number of parameters that define a spectrum allocation, the sorting problem may be very complicated. For example, there are a large number of categories that vary according to, among other things, waveforms and time. One ambitious aspect of the SC2 program is its desire to examine whether the rules of a highly decentralized spectrum management system can endogenously identify, via the collaboration game, a set of parameters which leads to a considerable level of spectrum sharing.

[25] The collaboration challenge may best be described as a non-zero sum game in that it has both cooperative and competitive elements. It is well understood that two players may have difficulty finding a suitable (e.g., Nash) equilibrium in such games. Further, the challenge players have in finding a suitable equilibrium rises dramatically with increases in the number of players.

[26] In markets where there are buyers and sellers, prices serve to coordinate the interests of buyers and sellers by sorting out for buyers the set of sellers that are most interested in selling to buyers, and by sorting out for sellers the set of buyers most willing to buy from them.

[27] Excess demand occurs when demand exceeds supply at zero price.

[28] It may be possible to examine the performance properties of a one sided-market where spectrum access is shared on a real-time basis, versus a two-sided market in which buyers have the ability to sell acquired spectrum. Under a one-sided market, the market consists of only users. Sellers play no role in the market since spectrum becomes available for use as soon as a user has exhausted its use rights.

10 Appendix D: Data Management and AI/ML Systems

AI/ML Data Systems

Like other Federal agencies and departments, the FCC manages important databases (e.g., service operator outage data), several of which can be made more valuable to private users and society through the application of AI/ML techniques.[1] The resulting analytics are likely to have a significant impact on all of the FCC’s high level strategic goals of promoting an innovative agenda, serving the under-served, protecting consumers, and advancing FCC processes. While data is very important, data simply represents one component of a broader AI/ML Data system. Figure D.1 (derived from [7]) below describes the various components of such a system. The size of each box roughly corresponds to the code/infrastructure that is needed to complete the identified task. There are several notable features of an AI/ML Data system. First, the process of collecting and verifying data, which includes such things as data collection, curation, labeling and storage, are not only two essential features of an ML system, but they also require a considerable amount of resources for their completion.[2] Second, only a small fraction of the overall code is actually attributed to the *ML Algorithms*.

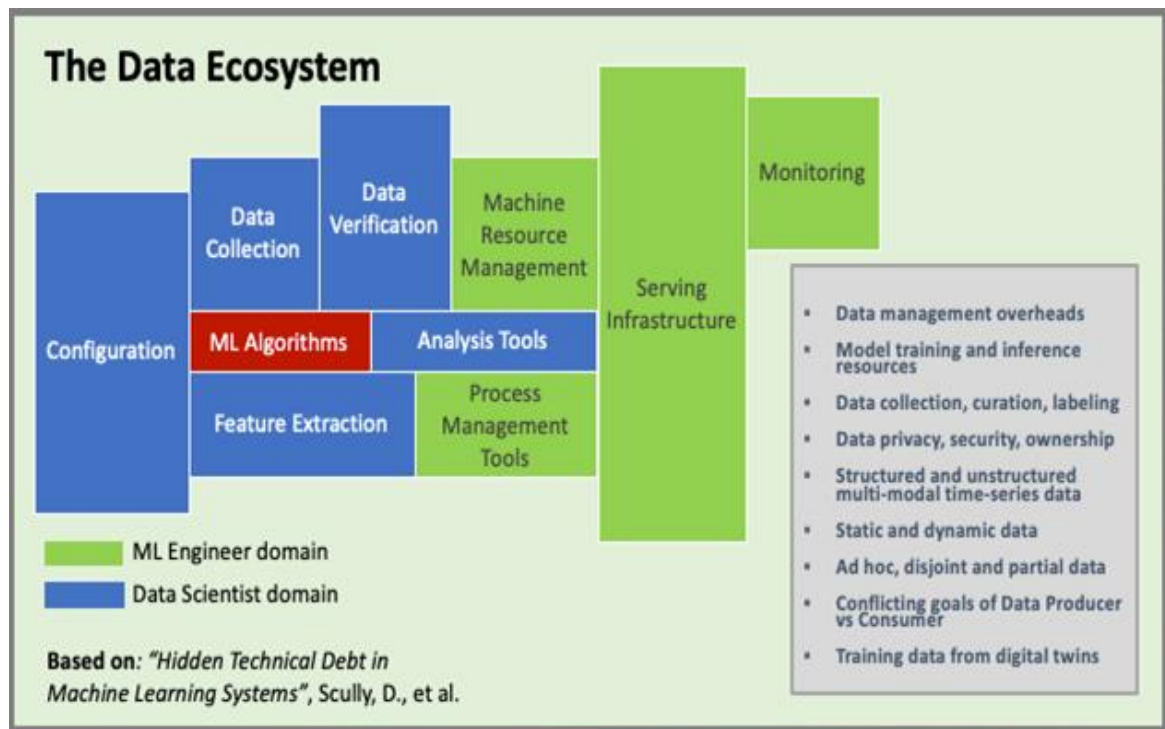


Figure D.1

Figure D.2 presents a flowchart that describes the sequence of steps involved with using ML to train an algorithm/model, as well as the amount of effort associated with each step.

Consistent with the information provided in Figure 1, a significant amount of effort and resource is directed towards developing and curating data. Figure 2 also displays the recursive nature of the training and learning exercise, including the measuring and assessing the accuracy of the predictions made by the algorithm.

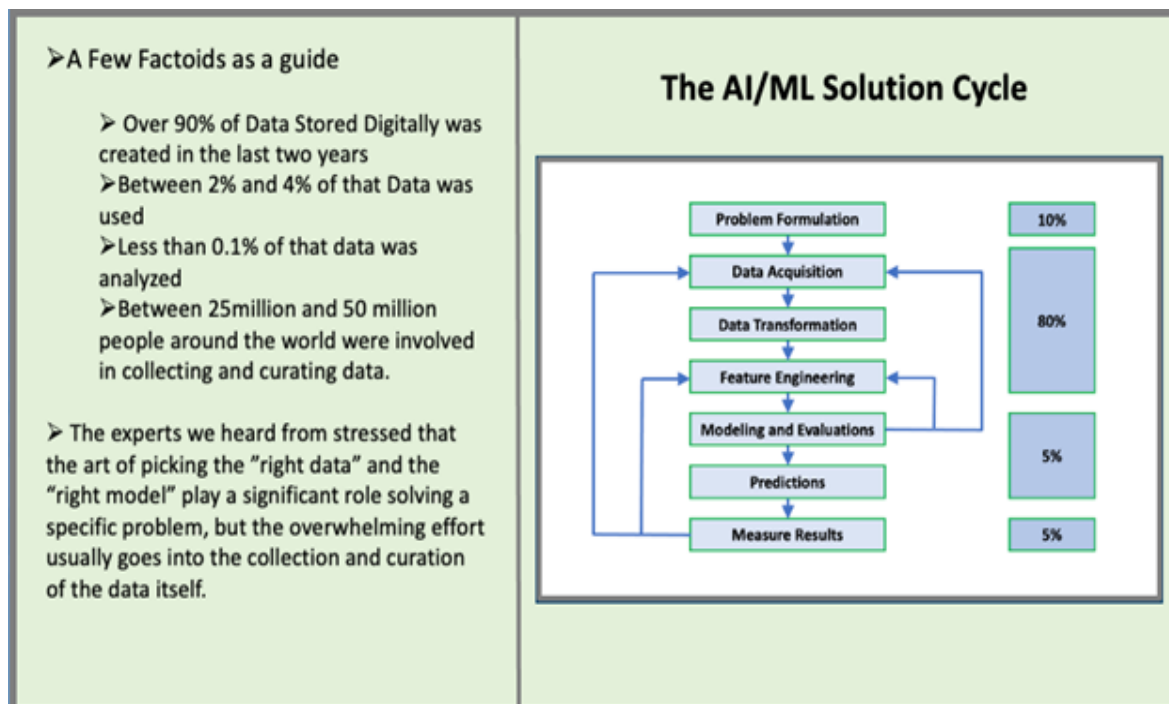


Figure D.2

The AIWG firmly believes that the benefits that flow from the FCC’s decision to develop and maintain well-curated data sets far exceed the cost incurred from doing so.[3] In addition to improving policy decisions, well-curated FCC, machine-readable data will be used by industry and academic researchers to build a variety of AI/ML models the result of which would promote competition and spur innovation in the telecommunications sector. For example, the existence of well-curated data introduces an intriguing new possibility. Assume for the moment that a standardized data set involving the performance properties of a set of telecommunications network components existed. AI/ML techniques can be employed to develop, based upon that data, a “benchmark” (i.e., point of reference) regarding the performance of each of those components over a wide variety of operational environments. A comparison of that benchmark with the actual performance of a firm’s chosen network component would allow the operator to make more informed purchase decisions regarding which inputs to include in its network. The operator’s ability to make more informed purchase decisions regarding components will spur competition and innovation in the network component market and improve the customer experience. “AI/ML performance benchmarking” can be applied to other areas of the telecommunications sector

including, for example, the performance of autonomous vehicles. In this case, AI/ML benchmarking can lead to improvements in the operational performance and safety of such vehicles.[4] In both cases, AI/ML benchmarking can provide important benefits for customers, the economy, and society generally. However, such benefits are contingent on the development of a fair and useful benchmark, defined as one that is based on a large, valid data set capable of allowing for the proper measurement and assessment of ML-based algorithms, hardware, software, and services. A widely accepted benchmark suite will benefit the entire telecommunications service and equipment community, including researchers, developers, hardware manufacturers, builders of machine learning frameworks, cloud service providers, application providers, and end users. One commonly overlooked aspect of AI/ML benchmarking activities involves the positive benefit that participating firms each enjoy by sharing data. This is important because, in many instances, owners of well-curated data have too little incentive to share their data with other parties. In such instances, the benefits to society from the additional knowledge that ML and data can produce are foregone.[5]

Strategies for Promoting the Sharing of Privately-held Data

In many instances, the absence of well-curated data is not the result of technical or financial matters, but rather because of the private-welfare maximizing interests of their owner. In many cases such owners can have “too great” of an incentive to not share their data from society’s perspective in that the gains that society obtains from sharing the data exceed the cost the data owner incurs from that sharing. The resulting effect is a diminished level of innovation in the U.S. economy broadly, and a reduction in consumer and producer welfare in the telecommunications sector. Importantly, the FCC can play an important role in correcting the misalignment of interests between society and the owners of the data. First, the FCC could encourage data owners to adopt new and innovative privacy preserving techniques to their data.[6] Second, the FCC could, by organizing and participating in meetings involving interested parties, facilitate solutions to the various hurdles that exist involving the sharing of data held by the owner of the data. Third, the FCC could begin to explore the feasibility of promoting the development of a data exchange. By enabling an owner of data to better monetize its data, a well-designed exchange may go a long way to increase the incentive owners of data have to share their data. Fourth, the FCC could set an example of the importance of data, and the application of modern analytics to that data, by making every effort to follow through on the AIWG’s recommendation for the preparation and publication of a “Request for Information” (RFI) that solicits information from the private sector on the various new and modern tools that can be brought together for the purpose of building and maintaining accurate wireless broadband maps.

[1] Appendix B provides a summary of some of the areas in which the Federal Government provides research funds for investigations in the application of AI/ML techniques to important resource allocation problems. The summary identifies numerous instances in which Federal entities are developing and managing important databases for use by AI/ML techniques. For information about data sets managed by the FCC, see <https://www.fcc.gov/general/broadband-deployment-data-fcc-form-477>. See also <https://www.fcc.gov/licensing-databases/search-fcc-databases>.

[2] “Curation” is the organization, integration and presentation of data from various sources in a manner that preserves its usefulness and reuse over time.

[3] The AIWG believes that the benefits from converting many of the FCC’s data sets into machine readable, well curated data sets far exceeds the costs associated with completing that transformation. In the event that the FCC believes there are better uses of its funds, the AI WG implores the FCC to explore the opportunity to share the cost of that transformation with other Federal Agencies and Departments.

[4] There are numerous other areas where the establishment of a standardized AI/ML benchmark would yield significant benefits. Those areas include the development of a benchmark involving RF interference involving 5G deployment, including mmWave 28 GHz and beyond. Another area involves a benchmark RF interference for existing LTE frequencies.

[5] In an effort to improve the rate of innovation in machine learning and to broaden its public access, some firms have come together to form an open consortium that offers individuals and firms access to various ML tools and databases. One of the data sets that is provided on an open access basis through the consortium is People’s Speech, the world’s largest public speech-to-text data. See <https://mlperf.org>.

[6] See presentation by Rafail Ostrovsky (UCLA/Stealth Software) “Stewardship of Private Data with Cryptography.”

[7] Scully, D., et. al., “Hidden Technical Debt in Machine Learning Systems,” NeurIPS 2015.

11 Appendix E: Safe Use of AI

Table E.1 below includes a list of the attributes of AI safety that need consideration in the design and deployment of AI/ML models in the context of applications, business processes, network planning and control. Note that we do not include privacy in this list; while it is an important consideration, it is not unique to AI/ML.

Fairness¹	<ul style="list-style-type: none"> • At individual, community and societal levels
Transparency²	<ul style="list-style-type: none"> • Of both models and training Data
(Non) Vulnerability³	<ul style="list-style-type: none"> • To adversarial attack
Accountability & Explainability⁴ to Humans	<ul style="list-style-type: none"> • Guided by the judgment of domain experts • Human compatible AI • Avoids harm to users (e.g., driving addiction, bullying, changing human behavior)
Robustness	<ul style="list-style-type: none"> • Prescribed operating range • Predictable response to unanticipated inputs • Well characterized Interaction of multiple models⁵ • Known <i>blast radius</i>⁶ • Ensure skillful implementation • Verification⁷

Table E.1: Attributes of Safe AI

As mentioned earlier, there are multiple insertion points for AI/ML techniques in Telco/service provider networks and they operate at different timescales, ranging from

¹ According to a recent McKinsey report (When Governments turn to AI: Algorithms, trade-offs, and trust”), “The notion reflects an interest in bias-free decision making or, when protected classes of individuals are involved, in avoiding disparate impact to legally protected classes.”

² There should be a reasonable degree of transparency as to when AI/ML is being used, the types of algorithms in use and the nature and sources of the data being used to train those algorithms.

³ There are many types of attacks to be considered here. For example, can a given AI/ML algorithm be deliberately fooled into making adverse decisions? Can the data used to train a model be tainted so as to create an adverse result? Can a trained model be reverse engineered to reveal private information?

⁴ The results of AI/ML algorithms should be easily explained, ensuring one can comprehend how and why conclusions are formed. We heard from many guest speakers, and it is also shared in the McKinsey report, that “AI is most valuable when used to support, and not substitute for, human decision making. From the McKinsey report: “Algorithms should be explainable, especially in the public sector, where myriad stakeholders will review every step. And, to ensure successful adoption, public-sector users should pay particular attention to how AI solutions are deployed.” Adopters of AI/ML should also be mindful of and, where appropriate, accountable for potential harm to users.

⁵ These models may be geo-distributed and/or at different layers.

⁶ Systematic vulnerabilities and the second order impacts of individual components on other parts of the system/network should be understood and contained.

⁷ That the implementations perform as designed and anticipated.

milliseconds all the way to days/years. The graphic below depicts the various insertion points and the timescales over which AI/ML techniques may take effect.

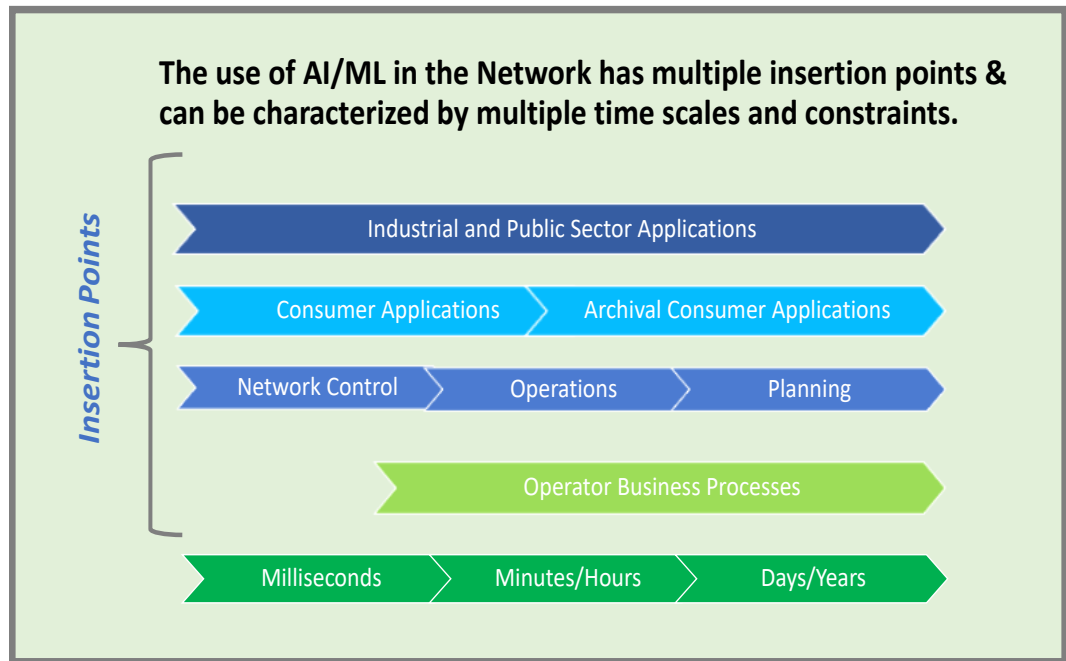


Figure E.1: Insertion points for AI/ML as characterized by applications and timescales

We next briefly discuss the above attributes of AI safety in the context of these insertion points.

(1) Consumer and Enterprise Applications:

The FCC will appropriately have concerns about Fairness, Transparency, Accountability and Vulnerability for AI/ML applications used by consumer services that are deployed on telco and service provider infrastructures. Some examples of these applications are Ad-insertion, content editing, real-time translation, application optimization, content analysis and indexing, record keeping, etc.. For example, the FCC may have accountability concerns related to the avoidance of harm to children.

(2) Operator Business Processes: Here the FCC will appropriately have concerns about Fairness, Transparency, Accountability and Vulnerability. Some examples of these business processes are: Customer service, data mining, billing, customer communications, outage reporting, etc. As an example, the FCC was recently involved in reviewing Robocalls and their impact.⁸

⁸ The growing widespread use of AI/ML may raise unique transparency concerns. For example, should the customer be informed about whether he/she is speaking with a human or an AI-powered computer during a service call?

(3) Network Planning:

The potential FCC concerns here include Fairness and Vulnerability. Some examples of AI/ML-powered network planning tasks that could be subject to unintended bias and/or adversarial attack are: Base station site selection, backhaul capacity provisioning,, AI/ML-powered service area decisions, etc.

(4) Network Control:

The potential FCC concerns here include Fairness, Explainability, Accountability, Vulnerability and Robustness. *Online* AI/ML, embedded deeply within the network and operating at *faster than human* time scales, has the potential to create significant challenges with respect to the vulnerability and robustness of the network – and the ability of human operators to oversee its operation.

Depicted below is an illustrative framework for the assessment of the robustness of AI/ML models that implement network control functionality. The illustration includes examples of specific questions related to the safe AI attributes that will need to be asked to assess the safety of the overall system.

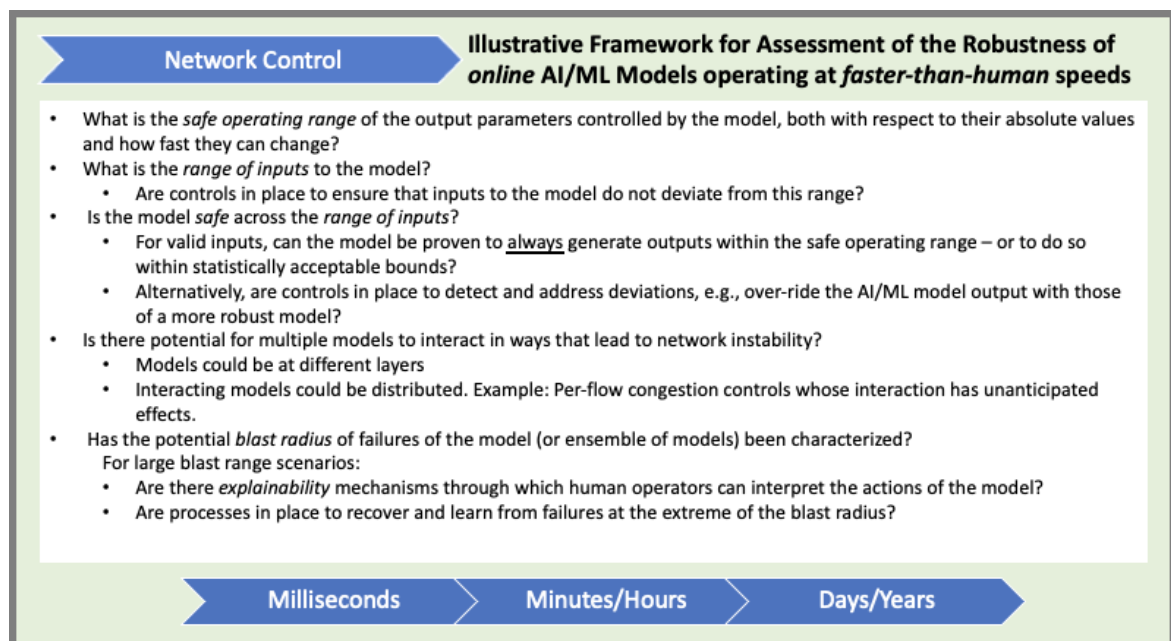


Figure E.2: Considerations for Assessment of Robustness for Network Control Functions

Instead of having each operator create its own framework, we suggest that the FCC encourage the creation of engineering consortia to develop one or more common assessment frameworks along the above lines. This would have the added benefits of creating a shared vocabulary, widely agreed criteria and a talent pool of skilled assessors. Figure E.3 below illustrates similar cases in which frameworks have been created in the past.

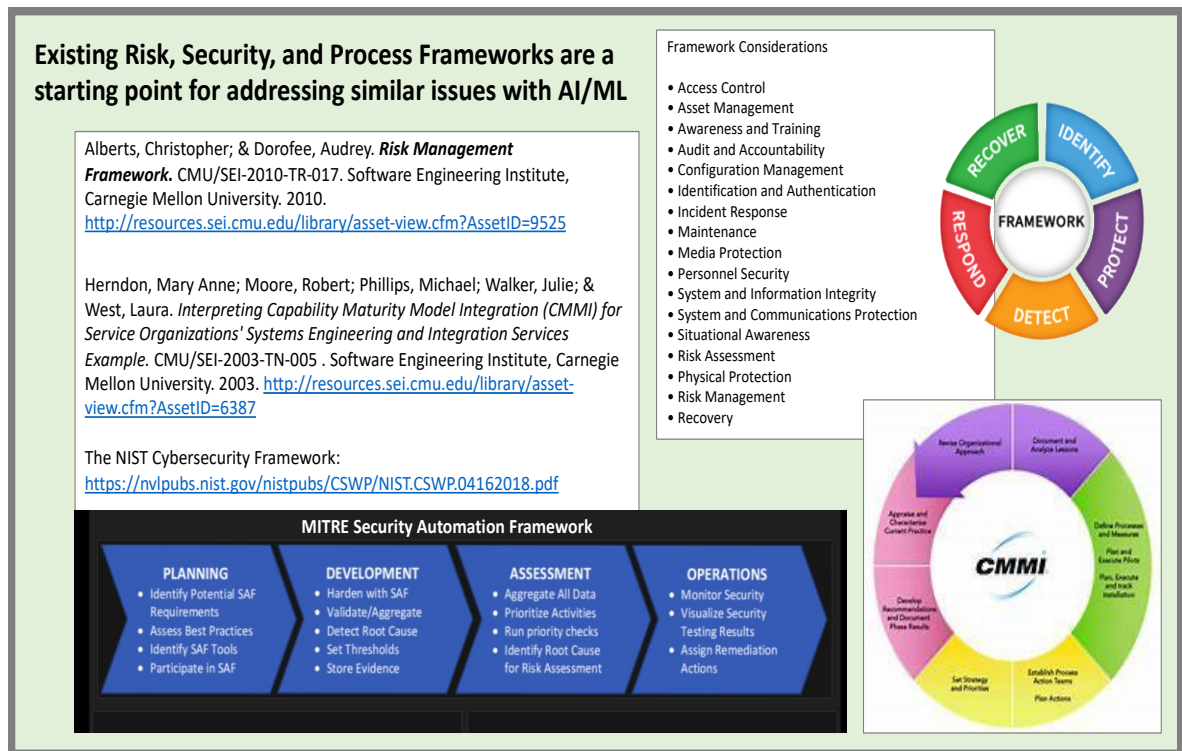


Figure E.3: Examples of existing assesment Frameworks

Expert Input

We had the pleasure of learning from two robust sessions with Stuart Russell (Professor of Computer Science at UC Berkeley, and author of several important books on this subject including *Human Compatible: Artificial Intelligence and the Problem of Control*), during which he spoke with us of the shortcomings of Artificial Intelligence See below Slide 1 in which “Google ponders the shortcomings of machine learning” which proposes we should seek a path where computers generalize more broadly about the world.” While the upside is “an enormous increase in the capabilities of civilization,” there are downsides such as impersonation, bias, robotics replacement of jobs, and the unimaginable outcomes, in cases where AI “outstrips our feeble powers.” Russell suggests AI be designed to be explicitly uncertain about human preferences, and thus need to turn to humans for judgment (see Slide 2).

Our McKinsey report, and other guest speakers in addition to Russell, all spoke of the better outcomes derived from a “human plus machine” approach. As an example, AI can reach too rigid a conclusion, especially in nuanced situations. For example, an “algorithm might deny a loan for an applicant with a credit score of 728 when the cutoff is 730.”

Russell further suggests that AI algorithms are best designed to “negotiate data exchange ... and thus our privacy.”

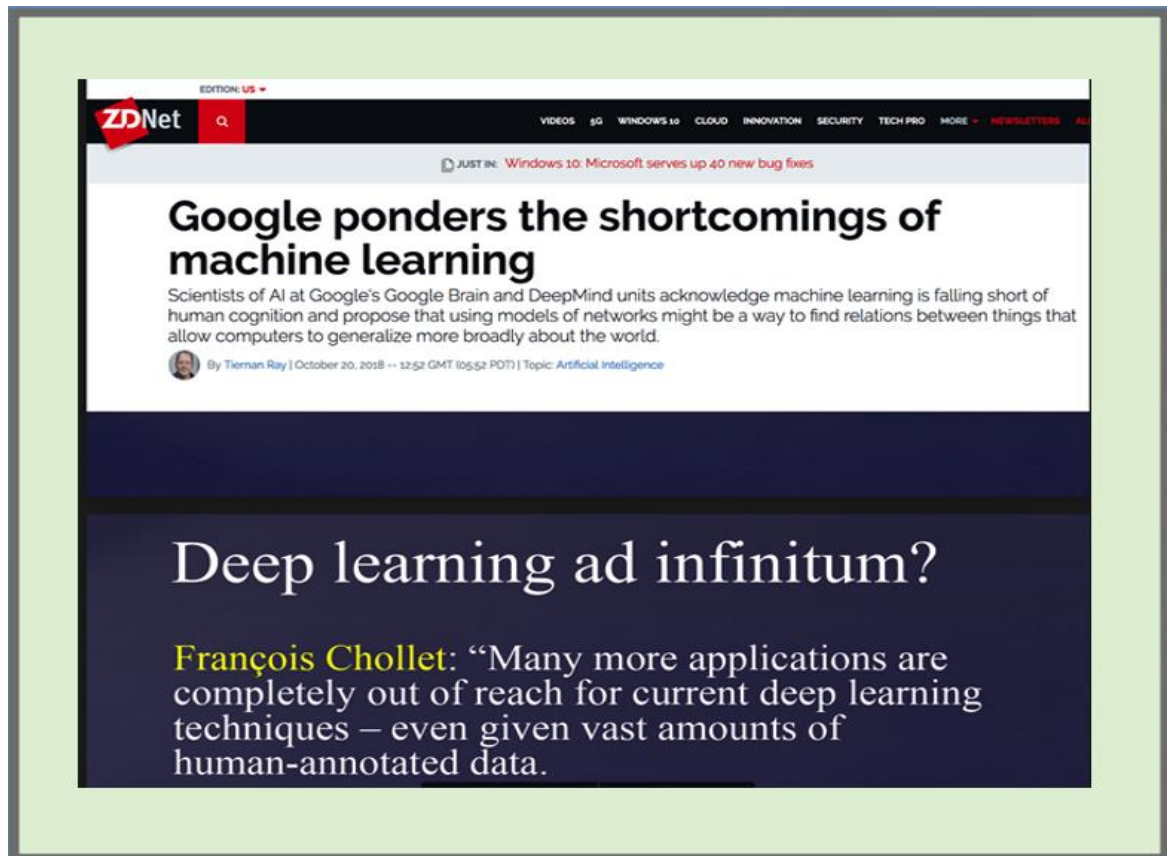


Figure E.4: Shortcomings of AI/ML Approaches

12 Appendix F: Uses of AI/ML in the Creation of a Broadband Map

Broadband Map Capabilities

The National Broadband Availability Map (NBAM) and the Broadband Deployment Accuracy and Technological Availability (Broadband DATA) ACT as funded by Congress directed NTIA with the coordination of FCC to acquire and utilize data from available third-party datasets to construct a map. The Broadband DATA ACT enhances the FCC ability to fix the existing broadband mapping inputs. The updates will require fixed and mobile providers to submit standardized broadband availability maps and for the FCC to develop a common dataset of homes and businesses to measure coverage against.

The development of a Broadband Map has many aspects to it and is a major undertaking. The project involves capturing the drivers and constraints, an understanding of the fleshing out of the requirements at a detail level, the concepts of operation for how it will be used and by whom and provisions for eventual deployment and sustainment. There are multiple approaches possible. The development entail, inclusion of a systematic analysis of the trade-offs for how the National Broadband Map will be built, what technologies will be involved, and the concrete plan for how the work will be accomplished.

Organizational Skill Set Needed for Broadband Map

The implementation of a revised Broadband Map can be met through a competitively offered solution. Organizations with experience in technologies, developing and integrating multiple approaches, operational and managerial abilities for projects of similar scope and content should be considered.

In presentations to the AIWG several experts indicate that significant applications of AI/ML technologies are already being used with success for network planning. AI/ML methods and technology uses are illustrated in Figure F.1. The requirements for such planning parallels many of the requirements for the Broadband Map as do the requirements for data collection. As with planning systems, the development of the National Broadband Map should be approached holistically and recognized the breadth of technologies, integration methods, and operational considerations that must be brought to bear.

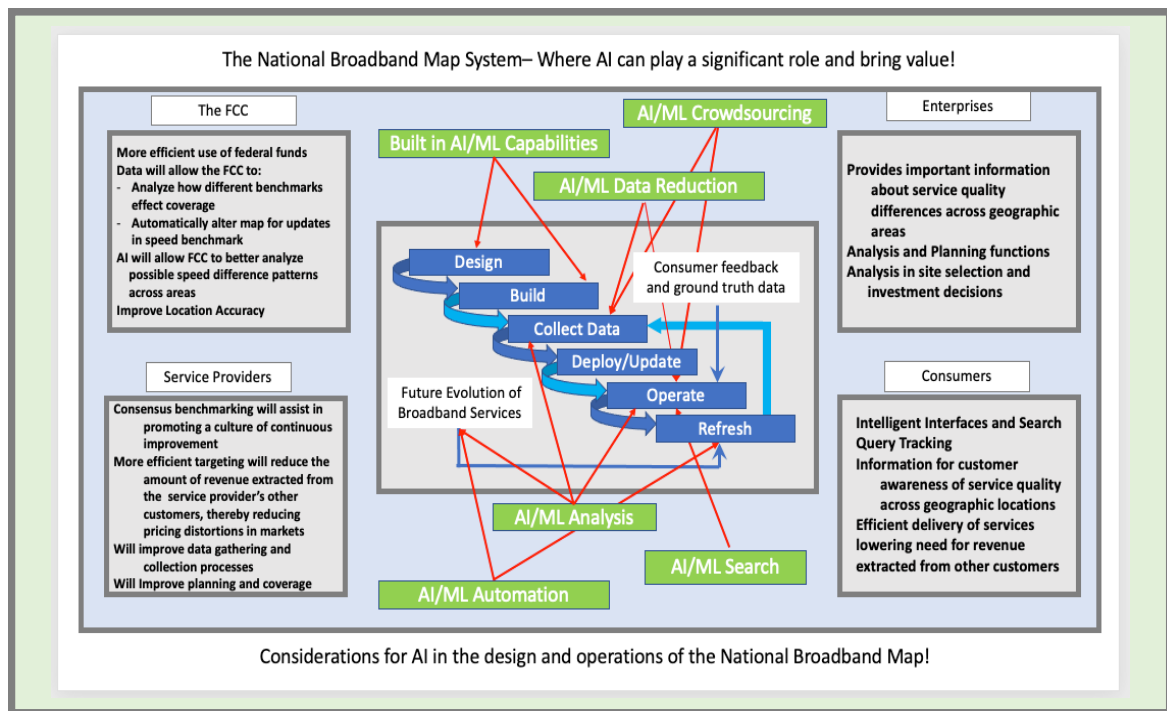


Figure F.1: Areas of benefits for the use of AI in the National Broadband Map

Implementation of Broadband Mapping

The development of a Broadband Map is a large project with many layers of data sets and considerations. A constructive step by step plan with timelines, milestones, a budget and an operational plan are required.

The Broadband Mapping project requires the collection of data, analysis of data, how the data will be used and by whom. Implementation of a Broadband Mapping plan requires the establishment of an accurate base map, receiving standards base data from existing broadband providers, developing methods for effortless updates and the ability for adjustable analysis of data with comprehensive results.

The overall project will also require integration skills, have a significant number of components and elements, field teams, and including data from many different sources in which AI/ML will play an important role.

There is significant capability available commercially and within academic institutions that a well-crafted Request for Information (RFI) will be informative to the FCC. The use of an RFI will provide a process to allow for the discussion of how AI techniques can be used in the development of a Broadband Map. The RFI path and how to set up the procurement process for the Broadband Map are discussed in the sections below.

RFI Path

The RFI Path is a resource tool for procurement of material and services that are outside the norm of daily purchase activities. The RFI has general uses, which allow for a broad range of preliminary requirements, project drivers and constraints to be openly discussed to enhance the knowledge base of the FCC prior to the selection process. The RFI process will also open the pool of responders allowing the FCC to receive a wider range of ideas or methods to meet the objectives being introduced and give experts outside of normal channels opportunity to present new processes which may enhance the overall scope of the broadband project.

The RFI will allow for the development of preliminary requirements, drivers and constraints in the process. Examples for the Broadband Mapping RFI include how Congressional mandates can be met, how the collection of information can be used by the FCC for the broadband map and how it can be used to meet FCC's NPRMs. The RFI will also allow for input by other government entities, service providers, commercial users and consumers to strengthen the need for select items or provide feedback on how a particular process will affect current private operations.

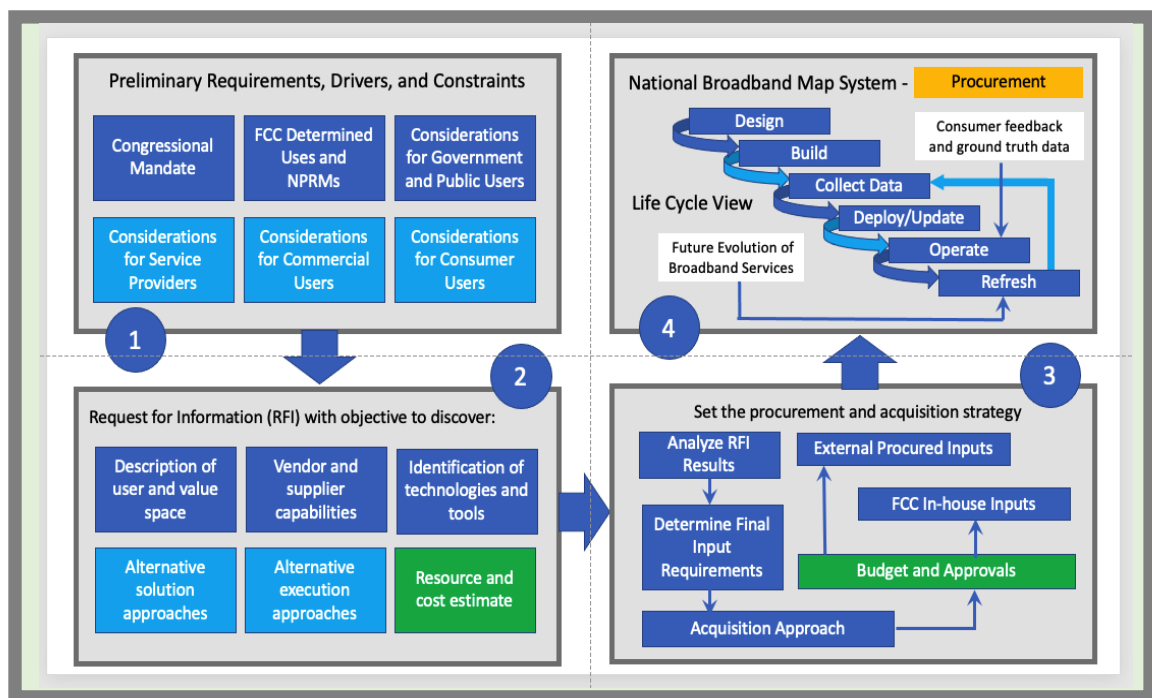


Figure F.2: A procurement process for acquisition of National Broadband Map Capabilities

A key part of the RFI process allows for an open gathering of information and objective discovery. The core components of the RFI allows for the description of user, added value space and the understanding of the vendor and supplier capabilities. The RFI process will

also give the FCC the ability to identify technologies and tools, alternative solutions and approaches, and the presentation of alternative execution approaches. With RFI responses, the FCC would receive information on resource and cost estimates.

The RFI process will allow the FCC to meet the broadband mapping requirement by identifying AI technology alternatives, integration approaches, development methods, deployment and operation approaches for provisioning as well as estimates of resources needed for each portion of the NBM's life cycle.

The use of the RFI will allow the FCC to receive a deeper understanding on data gathering techniques and how the data can serve the long term needs of the FCC by developing its AI/ML approaches.

How to set up the Procurement Process

The procurement process using an RFI adds a few steps but ensures an additional balance to lightly known topics/technologies or to large procurement projects. The RFI process includes the development of general questions, basic end requirements, intermediate steps needed in the process and payment structure. Once the RFI is received the FCC can analyze the information received. The RFI data can be used with FCC internal staff inputs and external procurement inputs to develop the final specifications for procurement.

13 Appendix G: Pilot Project Descriptions

During our deliberations the FCC TAC AI Working Group identified several topics on which targeted projects with near term focus could not only benefit the FCC in developing a better appreciation of the technical and practical aspects of AI solutions, but also progress FCC's longer term technical agenda towards meeting its strategic objectives. Several of these topics and associated benefits to the FCC were highlighted in Table Y. We also provide a more detailed exposition of topics that the AIWG deemed of importance as amenable for making progress in the near-term.

13.1 Pilot on Testing Lab

Description: FCC Laboratory facilities can benefit from a realistic network emulation capability which is properly benchmarked and calibrated and allows for testing and independent benchmarking of various factors that impact spectrum use and wireless networking performance. The emulator can also be very valuable for generating ground truth data that can benchmark and drive development of AI/ML solutions.

The emulation capability should factor realistic commercial deployments of current wide area well as local area communication networks. In particular it should emphasize testing and benchmarking in the higher spectral band regime (28-70+ GHz, for example), wherein the current emulation technology is not well-developed and propagation characteristics are not well-understood. Other usages could include evaluation of spectrum sharing policies, impact of spectrum assignments and related interference effects, benchmarking of propagation characteristics, to name a few.

A candidate emulator that can serve as a baseline for more sophisticated testing capabilities is the Colosseum Emulator developed as part of the DARPA SC2 challenge and now maintained by Northeastern university. The Colosseum will need to be re-architected and extended substantially to usefully service the requirements of commercial grade wireless telecommunication networks. A detailed description of existing capabilities at Northeastern University are available in an FCC TAC AIWG presentation by the group [Melodia, Gosain]. Noteworthy is the capability to deploy and test solutions of Open-RAN's "RAN Intelligent Controller (RIC)," functionality on the Colosseum., which will be very useful in developing and benchmarking a range of AI solutions for the next generation radio networks spanning from mobile devices, multi-antenna schemes, end-to-end network control including RAN slicing, to name a few. Additionally, there is potential to interface the Colosseum with NSF PAWR (Platforms for Wireless Research) testbeds which can be used for real-world testing of solutions developed in the emulation environment (e.g. COSMOS being deployed in the New York area).

Time frame. Working with Northeastern University, a basic emulation capability using the Colosseum framework may be set up in the first year. The FCC should partner with third-

party vendors and other government agencies to derive requirements and gradually expanded the capability in the next 3-5 years.

Specific Action for FCC. The FCC, potentially working with a third-party vendor and in collaboration with other government agencies (EX: NTIA/ITS, NIST), needs to develop new requirements for the Colosseum that are better suited to address commercial telecommunication networks. The emulator should be capable of emulating spectrum propagation and interference characteristics relevant for communicating with massive number of multiple antennas, in particular addressing higher band operation. It should also be capable of modeling realistic traffic and applications encountered in current networks as well as look towards emerging applications. The emulator should be flexible in adapting its configuration to model different deployments (urban, suburban, rural, etc.). This emulator should be benchmarked and calibrated against existing measurement data already available at the FCC ([ref?]). The FCC may also consider driving new measurement campaigns for more accurate calibration, potentially partnering with agencies such as NTIA/NIST/NSF.

The FCC may not only use the emulator and the datasets for its own use, but it can also consider opening the emulation facilities and datasets for research and development purposes, perhaps again collaborating with other agencies such as NIST, NTIA, and/or NSF.

Rationale. Having a sophisticated emulation and testing capability in the public domain is necessary, as commercial vendors and operators treat this capability as their competitive advantage. This capability will allow the FCC to use objective and quantitative basis for formulating policy, adjudicating and enforcing decisions related to use of spectrum, etc.

Several presentations to 2020 AIWG highlighted that the FCC is under-invested in areas that are critical to enhancing next generation communication networks. As an example, it has been pointed out that wireless propagation models currently in use are inadequate for developing properly benchmarked solutions for next wave of smart antennas solutions, operation in higher spectral bands, and propagation in cluttered environments, etc. AI offers significant potential to assist and develop enhanced models, which can be benchmarked under FCC's guidance and made available broadly. Strengthening FCC's ability to invest in both technology expertise and infrastructure capabilities will be important to spur innovation and to drive US leadership in developing next generation of networking technologies.

Benefit. A realistic in-house network emulation and testing capability is expected to have a multi-fold benefits for the FCC. It will allow the FCC to accurately replicate and examine network conditions with high confidence. For example, it will allow FCC to better quantify new technologies for spectrum sharing and reuse and improve decision making with regards to spectrum policy. A properly calibrated emulator with field data can also be used to generate data sets that drive AI/ML innovation for next generation telecom networks, filling

a critical gap in this area. FCC can also make this testing facility available for non-commercial research and experimentation in partnership with agencies such as NIST or NSF, which could spur more research and broader innovation.

Mapping to govt bureau (s) The Office of Engineering and Technology (OET), The Wireless Telecommunication Bureau (WTB), The Enforcement Bureau (EB), The Office of Economics and Analytics (OEA) and, The Public Safety and Homeland Security Bureau (PSHSB).

Mapping to strategic objectives. A state-of-the-art testing lab fundamentally impacts all the objectives laid out in FCC's strategic plan:

Closing the Digital Divide: The testing lab when combined with focused efforts in developing new propagation models, can assess and promote new methodologies for sharing scarce spectral resources. Potentially new spectrum can be made available under sharing arrangement with significantly lower cost, thus providing access to underserved communities, bridging the digital divide. The FCC will have the tools it can use to help close this digital divide, bring down the cost of deploying broadband, and create incentives for providers to connect consumers in hard-to serve areas."

Promoting Innovation. The capability of generating datasets that are benchmarked will catalyze innovation in AI for telecom networks. Several presentations have noted that availability of benchmarks such as ImageNet, contributed to success of AI innovation for applications such as image recognition and object detection. Similar benchmarks available for public use are considered key to success of AI in the wireless/telecom domain. In addition to data sets, FCC's testing and emulation facility will also help to benchmark new innovations going beyond AI solutions. In particular, "it will ensure that the FCC's actions and regulations reflect the realities of the current marketplace, promote entrepreneurship, expand economic opportunity, and remove barriers to entry and investment."

Protecting Consumers and Public Safety. The testing facility can be used to evaluate the extent of harmful interference caused by unlawful, or inappropriate use of spectrum.

Reforming the FCC's Processes: The testing lab significantly enhances FCC ability to make data driven, quantitative decisions. For example, interference issues arising from spectrum sharing may be addressed in an objective manner. Specifically, it will "modernize and streamline the FCC's operations and programs to improve decision-making, build consensus, reduce regulatory burdens, and simplify the public's interactions with the Commission."

Impact to the US. State of the art testing and data generation facility will foster innovation especially as it relates to AI/ML technologies, which will enhance telecommunication networks and promote fair use of spectrum and networks. Overall, it will be important in driving US leadership in developing next generation of telecommunication networks.

13.2 Pilot Project to Improve Innovation Using AI/ML Benchmarking

Background. A key priority for the FCC is to foster a competitive, dynamic, and innovative market for communications services through policies that promote the introduction of new technologies and services. We will ensure that the FCC's actions and regulations reflect the realities of the current marketplace, promote entrepreneurship, expand economic opportunity, and remove barriers to entry and investment.

Rationale: Currently no standardized AI/ML Benchmarking exists for:

- RF Spectrum Allocation analysis specific for 5G System deployment
- mmWave 28 GHz and beyond
- Interference for existing LTE frequencies
- RF Propagation Model development

Currently no standardized data sets available for competition and innovation.

Specific Action for FCC. Improving Innovation through AI/ML Benchmarking

- FCC to create an RF spectrum AI/ML Benchmarking team (RF AI Team)
- Proposed RF AI Team charter
- Audit existing FCC data sets to determine relevance to 5G frequency deployments
- Publish report on existing data sets
- Work with industry and academia to propose new AI/ML models:
- Create criteria for useful and valuable predictions from AI/ML models
- Improve modeling of RF Spectrum Interference
- Improve propagation models
- Create a set of benchmarks to measure model performance
- Work with industry and academia for standardized performance model benchmarks
- Create uniform Compliance Test Specification for model evaluation

Existing AI/ML Data Sharing Examples and Benefits

- ML-Commons and Mlperf.org define AI/ML performance benchmarks
- <https://mlperf.org/>
- <https://mlperf.org/press>

- Require VAST amounts of training set data
- BERT: Bi-directional Encoder Representation from Transformers (BERT):
SQuAD 1.1 data set
- 3D U-Net: BraTS 2019 dataset for brain tumor segmentation
- RNN-T: Recurrent Neural Network Transducer: LibriSpeech dataset

13.3 Pilot Project on AI to Aid Enforcement

Background: The FCC’s spectrum related enforcement activities can be divided into two categories – ex ante and ex post. Interference is unwanted (RF) energy that disrupts the reception of desired information. Ex ante activities relate to actions like consumer education or requiring product labelling that are intended to deter interference before it occurs. The latter relates to enforcement activities that occur starting when the interference manifests itself. It is widely agreed that there are six steps that need to be taken in ex ante interference enforcement: (1) detect, (2) classify/identify, (3) locate, (4) report, (5) mitigate, and (6) remediate.[1] AI/ML techniques could be useful in each of the six steps, especially when the interference is associated with deliberate, malicious jamming, spoofing, and, potentially, “sniffing”/unauthorized interception. AI/ML techniques can be used to detect unwanted energy that is disrupting a communications link and to classify/identify it. To take a simple example, such techniques could be used to not only detect the presence of unwanted energy on a link, but also to classify or identify it as, say, the harmonic of an FM radio station. Continuing, the techniques could be used to geographically locate the FM radio station and to automatically report on the incident. By collecting the details of the FM radio station’s signal using RF Fingerprinting techniques, essentially irrefutable evidence could be assembled and used to seize and thereby stop (that is mitigate) the operation of an unlicensed pirate station and to prosecute (that is, remediate) the individual or group responsible for the interference.

Description:

Based upon the background provided above, the Rothwell and Figg presentation, and the desire to produce tangible results within one year, two potential pilot projects emerge:

1. As an outgrowth of previous work in the TAC, facilitate the establishment of an Interference Hunters Information Exchange (IHIE). Its role would be to allow the exchange of information among interference hunters on radio interference incidents, trends, and resolution and mitigation techniques for the purpose of protecting users of the increasingly valuable and congested radio spectrum environment against harmful interference of all types – incidental, unintentional, and both malicious and non-malicious intentional interference. Participants could include not only private sector entities (including wireless service providers and the manufacturers of spectrum monitoring equipment) but the FCC

and other agencies like the FAA which regularly collect information on interference incidents and their resolution. It could also include participants from neighboring countries including Canada and Mexico. It is recognized that, while the FCC may not always be able to release information on specific interference incidents because of on-going enforcement actions, even somewhat dated or general information could still be quite useful.

2. Establish an industry/government project to collect and curate information on the emission characteristics of systems and devices that are known to have caused harmful interference to existing wireless services. Such systems and devices would initially be identified by anecdotal information from various sources (e.g., professional and trade publications) and then expanded based upon the information collected as part of project 1., above. Information on the emission characteristics of frequently interfering devices could be augmented by the ex-ante information that the FCC collects in conjunction with its equipment authorization program. As part of that program, manufactures of certain electronic devices are required to identify the types of RF emissions their devices will generate using a list of standardized emission designators. Augmenting the interference information in this way would (a) allow wireless operators to more readily make design, implementation and operational choices to avoid or minimize interference and (b) provide the FCC with another ex post interference enforcement tool they could use to require a manufacturer to, for example, stop selling or even recall devices that are causing interference because of non-compliance.

Time Frame: If the two projects described immediately above were initiated early in CY 2021, they could be producing useful results by the end of the next TAC if not sooner.

Specific Action: Take the steps necessary to facilitate (1) the establishment and operation of IHIE and (2) the establishment and operation of industry/government project to collect and curate information on the emission characteristics of systems and devices that are known to have caused harmful interference to existing wireless services.

Rationale: The explosive growth in wireless devices and systems that operate in increasingly close proximity to one another in space, time and frequency, coupled with rapid technological, operational and business developments is changing the interference environment and putting increased pressure on traditional interference resolution and enforcement methods. For example, using Software Defined Radio techniques, the desired waveforms that can be generated by a transmitter to improve system performance become almost infinite in number and they can be changed in milliseconds or less. On the other hand, these same techniques can be used by perpetrators in a non-malicious or malicious fashion to more efficiently and effectively disrupt wireless devices and systems through jamming, spoofing and “sniffing”/unauthorized reception at the RF layer.[2] Such disruptions pose critical threats to our economic and social wellbeing and to national defense and homeland security.

Mapping to Relevant Bureaus: The two potential pilot projects proposed herein are especially relevant to the Enforcement Bureau (EB), the Office of Engineering and Technology (OET), the Public Safety and Homeland Security Bureau (PSHSB), the Wireless Telecommunications Bureau (WTB) and the International Bureau (IB).

Mapping to Strategic Objectives: The two potential pilot projects can contribute significantly to three of out of the four priorities emphasized in the FCC's 2018-2022 Strategic Plan. Namely, they can contribute significantly to Closing the Digital Divide, Protecting Consumers and Public Safety, and Reforming the FCC's Processes.[3] Rather than describe all of the possible contributions within each of these three priorities, selective examples in each are provided below.

- Strategic Goal 1: Closing the Digital Divide

One of the strategies that is called out for achieving Goal 1 is for the FCC to work closely with the National Telecommunications and Information Administration and regulators in neighboring countries to “identify and resolve instances of harmful interference on an international basis to avoid harmful interference in the future.” The first pilot project calling for the establishment of IHIE would contribute directly to successfully carrying out that strategy by facilitating the exchange of interference information while the second pilot project would contribute to it by providing information on the emission characteristics of systems and devices that are known to have caused harmful interference to existing wireless services.

- Strategic Goal 3: Protecting Consumers and Public Safety

The two pilot projects would contribute the most to achieving this Strategic Goal. For example, both projects would contribute significantly to Performance Goal 3.3.4 which calls for the FCC to “Provide situational awareness of communications systems; coordinate with industry and other Federal partners to facilitate communications network preparedness, response and restoration by working closely with local, state, Tribal governments during a crisis.” As a more specific example, having a catalog of the emission characteristics of systems and devices that are known to have caused harmful interference would be of a significant help in responding to interference caused outages.

- Strategic Goal 4: Reforming the FCC's Processes

The two pilot projects would also contribute to strategies singled out by the FCC to achieve this Strategic Goal. For example, they could contribute strongly to the expressed strategy of putting in place “processes that provide for timely introduction, upgrade, or replacement of technologies and identify ways to leverage and integrate technology to eliminate unnecessary redundancy and promote efficiency and effectiveness while maintaining continued adherence to a high level of information security standards.” Clearly the two pilot projects can, for example, lead to the use of AI/ML techniques that would promote the

effectiveness, efficiency and timeliness of the processes and tools that the FCC uses in carrying out its spectrum enforcement responsibilities.

Impact on the U.S.: The U.S., like the rest of the world, is becoming increasingly reliant upon wireless systems and devices that are inherently open to disruption. It is a simple fact that, in order for wireless-based systems to function, they must be open at the RF layer. Hence they are vulnerable to natural noise, e.g., radio “static” and cosmic noise, and manmade noise and interference of all types. The inherent vulnerability applies to not only communications networks, but to radio-based Positioning, Navigation and Timing (PNT) systems such as GPS/GNSS, radar systems, and remote sensing systems. The manmade interference can be unintentional or intentional and, in the case of the latter, malicious or non-malicious. Regardless of the intent, harmful interference can seriously disrupt wireless-based systems that are crucial to the economic and social well-being of the nation and, especially, to the safety of life and property and to the national defense and homeland security. The two pilot projects proposed herein would provide vital information that can be used (a) to increase the effectiveness, efficiency, and timeliness of the techniques used to detect, classify/identify, locate, report, mitigate and remediate harmful interference and (b) to provide data-sets that can be used to in the creation of automated tools and AI/ML techniques that hold the promise of substantial improvements over existing techniques.

[1] If both the device or system causing and receiving the interference are operating in accordance with the applicable rules governing their services, then an additional “deconfliction” step may be necessary. Spectrum enforcement is obviously easier when one party or the other is clearly at fault.

[2] In addition to jamming, spoofing, and sniffing, an additional type of disruption called “big data” is often identified. It involves the integration of multiple sources of RF information over time/frequency/space to determine identities or activities and the extraction therefrom of “pattern of life” analyses.

[3] The two proposed pilot projects are also relevant to the Promoting Innovation priority but less directly. For example, the FCC has promoted spectrum sharing as a way of improving the efficiency of spectrum usage. However, the willingness of entities to dynamically share their spectrum with others may well depend on the degree of confidence that they have that their spectrum will be protected from harmful interference from other systems and devices sharing the resource.

13.4 Pilot Project on Self Organizing Networks

Description: Self-organizing network (SON) solutions are continuing their evolution towards full automation. While traditional SON currently implements automation in targeted areas, there is still significant human involvement in network management and aspects of network control. As networks continue a steep rise in complexity and sheer number of nodes, with each new generation, near full automation is no longer optional. It is becoming a necessity and is important for supporting two objectives for the FCC, achieving a leadership role in Telecommunications, and supporting the innovation agenda.

This is also an area in which AI/ML is likely to play a dominant role. While the development of actual solutions is something that is most likely to be led by industry, the FCC can have a significant role in hastening the advances to mature the technology and in assuring that it is aligned with FCC policies.

From 2G to 3G to 4G to 5G and beyond, systems have become more and more complex. As an example, the operation of a typical country-wide network requires automation and AI-ML to manage it because the network could consist of:

- 100,000 to 1,000,000+ base stations
- 1,000,000+ configuration parameters
- 100+ critical events / hour

SON solutions will have to support:

- Self-configuration - Automated configurations tasks, like site creations
- Self-optimization - Optimized parameters and coverage usage for increased network quality and performance
- Self-healing - Automated fault resolution resulting in faster maintenance and reduced outage times reducing time and cost of system failures

Full automation will require the use of AI/ML. And this will require:

- Data (i.e., access to data from the network)
- On-line model re-training (esp. as the environment changes over time)

While we typically think of wireless networks as implementing SON, this type of automation may also be applied to fixed networks as well.

With respect to 5G, multiple logical networks on a shared physical infrastructure will be possible and will be optimized through network slicing. SON can be the solution for the 5G RAN NSSMF (Network Slice Subnet Management Function). In this capacity, SON will provide:

- Automatic radio slice life cycle management
- Radio network slice optimization
- Radio slice resource optimization

And in this capacity, SON will need to also ensure that it works well in the context of spectrum sharing. This may include the involvement of SON's ANR (i.e., automated neighbor relation) function.

In formulating a possible pilot, even a project of limited scope could resolve some of the challenges faced by researchers, equipment providers, and operators. It could also provide significant value to them and the FCC within a one to two year period.

Specifically, the FCC could:

1. Implement a testbed and the supporting infrastructure that enables and encourages the sharing some aspects of management and control data generated by the network for AI/ML models. This testbed should function as a mechanism for data coordination between data producers and data consumers as a neutral environment.
2. Experiment with integration of SON for network slicing on existing or upcoming spectrum sharing programs and partner with other relevant research and applied research agencies.
3. Look for ways to improve speed and efficiencies leveraging AI in addition to ensuring that SON plays well in spectrum sharing regimes.

With respect to a data management framework for the coordination of data between data consumers (e.g., AI/ML models), data collectors, and data sources, 3GPP TR 23.700-91 is defining such a framework. See Figure G.1 below. It would interwork with such entities and frameworks as the NWDAF – Network Data Analytics Function of 3GPP TS 23.288/29.520.

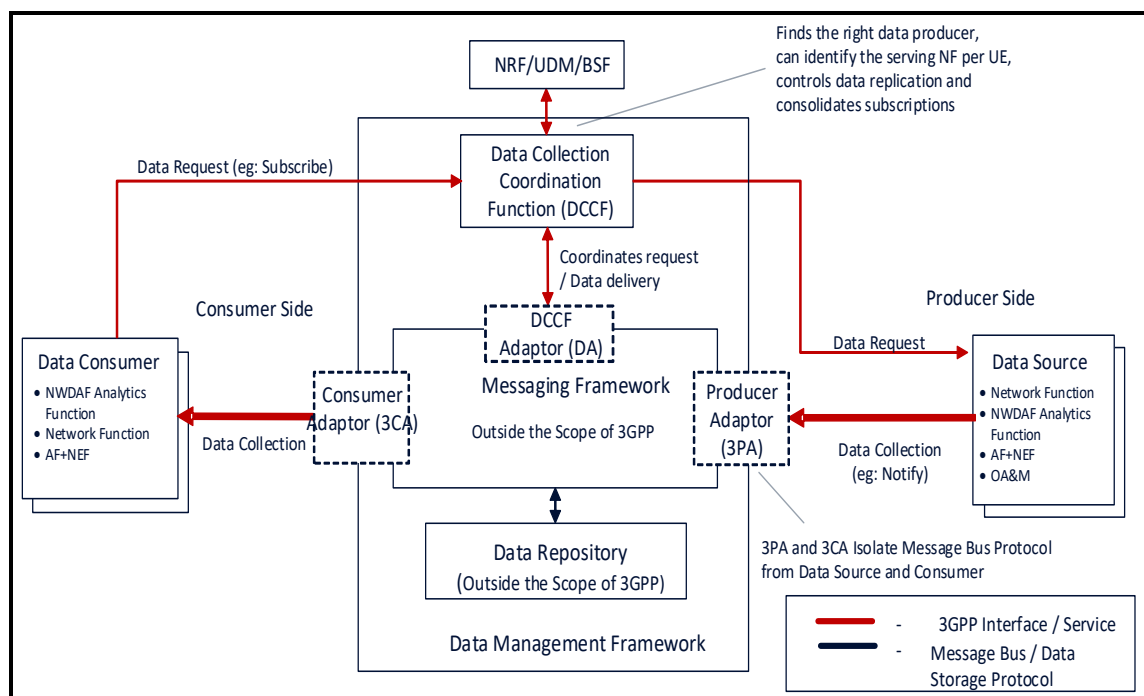


Figure G.1: The Data Analytics Functions as Defined by 3GPP

The FCC implementation of this proposal should be as closely aligned with standards as is possible with the understanding that in some cases this could lead to an early “Community” prototype.

The Benefits of the Pilot: The experimentation around the testbed, and it’s very availability will likely incentivize the sharing of data for AI/ML model development and learnings about the viability of SON approaches and solutions. With respect to the spectrum usage proposal,

this may impact spectrum sharing policies by leading to significantly improved efficiencies and new dynamic spectrum sharing regimes.

Overall, the effort will help the U.S. in achieving AI/ML world leadership for telecommunications. Currently, other nations appear poised to take the leadership mantle. Specifically, obtaining data for training is not as hard for vendors of some other countries. For the spectrum sharing part of this proposal, it should help to ensure efficient and fair spectrum usage amongst all players.

The effort should leverage the investment of other U.S. government agencies, especially the FCC should consider working/partnering with programs that could play a role in this pilot, including PAWR at the NSF (e.g., COSMOS), the NSF MLWiNS program, and the efforts at the NTIA in Boulder.

In summary, the results of such a pilot are also important and are likely to make a major contribution to for the future evolution of 5G, network slicing, automating SON, and spectrum sharing.

13.5 Pilot Project - Employing AI to Assess RF Interference and Promote 5G Open Radio Access Networks

Background:

The FCC's transmission power limits attempt to produce a spectrally efficient level of co-existence between and among the independent users of the RF environment. The level of spectral efficiency derived from those transmission power limits is often debated in the public record. Furthermore, the performance built into some receivers based upon existing market forces can itself lead to protracted disputes about theoretical thresholds of acceptable and unacceptable levels of tolerable interference.

Wireless networks are becoming increasingly complex, dynamic, and the demands and use of spectrum are increasingly non-deterministic and non-linear. Such rapid technological change raises further questions regarding the level of spectral efficiency of static transmission power and emission limits. As in the past, the FCC will be asked to intervene and weigh in on highly contested disputes regarding RF interference. Artificial intelligence (AI) and Machine Learning (ML) have the capability of substantially improving the level of spectral efficiency that exists between and among wireless service providers, by enabling heterogeneous radio systems to adapt to local RF environments to avoid interference without human intervention. This is because AI, modern statistical methods, and radio receivers and transmitters that can signal their coexistence to each other, can autonomously adjust signal power and reception paths. AI/ML-based solutions have the potential to differentiate between environments where there is a high likelihood of significant RF interference, versus environments where there is a low likelihood of such interference. Such an ability would greatly diminish the substantial cost associated with applying a set of

transmission power rules nationwide. Further, in environments where there is a high likelihood of RF interference, AL/ML-based solutions may be able to mitigate the level of that interference through instituting real-time dynamic interference avoidance measures.

A famous management guru once said, “you can’t manage what you don’t measure.” The thought applies with substantial force regarding the RF environment. While there are analytical models that capture, for instance, the non-linearity in radio transmitters and receivers, those models are typically incomplete and, thus, do not accurately describe both the day-to-day and extreme RF environment. The needed ingredient is data. Unfortunately, widespread, practical, public, real-world RF field measurement data of desired and undesired RF signals are typically unavailable. The absence of such data, combined with the incompleteness of the models, can lead to either a significant overestimation or an underestimation of RF interference and its effects on other radio services. For example, worst-case analytical models of RF propagation of fundamental and spurious/out-of-band emissions can lead to a significant overestimation of RF interference. On the other hand, no analytical analysis of certain RF phenomenon (e.g., passive intermodulation (PIM)) can lead to the underestimation of RF interference on radio services that are far removed (spectrally) from adjacent neighbors. The resulting inefficiencies hamper the FCC’s ability achieve its goal of fostering a competitive, dynamic, and innovative market for communications services.[1]

Furthermore, the general inability of radios of disparate radio services to intelligently “see” (from an RF standpoint) and avoid neighboring radios in their vicinity without the direction of human control or radio spectrum access system control, limits the deployment of some new radio technologies due to protracted policy and legal disputes before the Commission.

There are many instances in which real-world RF field measurements (data) regarding both desired and undesired RF signals are made. Such measurements are essential, for example, for the ability of multiple radios to share the use of common spectrum. These measurements are used, for example, by base stations in a licensed mobile broadband network for a variety of essential network control functions; for example, measurements to minimize co-channel interference from neighboring base stations that may use the same frequencies, measurements to enable base stations to ‘handover’ wireless links to maintain session connectivity of a mobile user devices that is physically moving throughout the network, and measurements to adaptively control the number of bits in a wireless connection to a mobile device to optimize the spectral efficiency (bps/Hz) of widely varying signal power and signal quality to mobile devices served by the base station. However, while the format and definition of such RF field measurements may be standardized (e.g., 3GPP standards), the measurements themselves are generally private and proprietary to the network operator and, thus, are not shared publicly.

Radio links of other types of radio services also perform RF field measurements and receivers act on the strength and/or quality of received signals to perform their radio service (e.g., broadcast receiver, radar receiver, microwave point-to-point receiver, land-mobile public safety receiver, etc.), and, generally, these measurements are proprietary and, thus, not publicly shared. One consequence of this is that radios of different radio service types generally cannot “see” each other specifically, in the sense of being able to detect and maneuver themselves out of the way of other heterogeneous radio services operating in the neighboring vicinity (spatially or spectrally).

TAC recommends that the Commission promote open RF measurements to advance the assessment of RF interference, improve system interoperability, and enable increasing coexistence of next generation wireless technologies.

- The TAC believes that open RF measurements can be a key enabler of Open-RAN and future RF network architectures, whereby, diverse heterogeneous radio services can signal and coordinate (e.g., for short durations of transmit power control to facilitate and mitigate interference between radio services), without requiring “human in the loop”, or even “Spectrum Access Systems (SAS) – in the loop”, to effectively avoid and mitigate interference between services such as satellite systems, radar systems, and aeronautical / terrestrial mobile broadband systems.
- TAC recommends that the Commission formally coordinate with NSF Platforms for Advanced Wireless Research (PAWR) to identify advanced RF measurement methods across heterogeneous radio services, services that have traditionally been siloed without the means of ‘cross service’ signaling to maximize the interoperable use of shared spectrum.
- TAC also recommends that the Commission formally coordinate with DARPA in programs such as Radio Frequency Machine Learning Systems (RFMLS), where capabilities such as RF Fingerprinting, RF Fingerprint Enhancement, Spectrum Awareness, and Autonomous RF System Configuration, could evolve to enable diverse radio system coexistence in the same or adjacent spectrum. Connecting these research projects with future FCC spectrum policies could enable Open-RAN and facilitate the application of AI / ML in 5G networks of the future.

TAC also recommends that a pilot program focused on the evolving Unmanned Aircraft Systems (UAS) use case could help inform Commission action on spectrum policies specifically focused on 5G Open Radio Access Networks (O-RAN),[2] Virtual Radio Access Networks (V-RAN), Artificial Intelligence (AI), and Machine Learning (ML).

- Why O-RAN – As mentioned above, RF measurements are generally held privately, not shared publicly, and are proprietary in conventional wireless

networks. In addition, functions such as radio resource management, which allocates radio channels dynamically on-demand, are proprietary and private. While protocols used for “handover” are more open in today’s networks, to enable interoperability between different network service providers, the universe of heterogeneous entities that autonomously interoperate at the radio layer between each other is limited. More importantly, FCC spectrum licensing for HF, VHF, UHF, radar, ADS-B, etc., are generally siloed and do not facilitate the spectrum needs of a single “UAS session” (i.e., cross-country, urban-rural-urban, command and control, detect-and-avoid, radar altimeters, used for flight from takeoff to landing). Nor does FCC spectrum licensing enable multiple “virtual” RANs needed by an evolving National Airspace System (NAS) that will not be solely operated for unmanned aircraft by the federal government, but by multiple UAS Service Providers.

- Why V-RAN (and O-RAN) – UAS spectrum needs have been allocated internationally in a single block of spectrum (5030-5091 MHz), which is limited in its propagation characteristics. Multiple frequency bands are needed. Just as manned aircraft use HF and VHF spectrum for long distance communication, radars, and transponders of various types, unmanned aerial vehicles (UAVs) have similar spectrum needs. While the dynamic bandwidth needs for machine-machine IoT UAVs may be less demanding than analog voice channels for some operations, the need for licensed protection rights (safety of life) is no less demanding. The current situation of dominantly using ‘unlicensed’ spectrum for UAS is unsatisfactory. A national seamless network of diverse licensed spectrum (low, mid, and high band) is needed. V-RAN brings the opportunity of licensing multiple nationwide slices of UAS spectrum rights, with single licenses that serve the HF, MF/LF, VHF/UHF, radar, needs of a single UAV or multiple classes of UAVs.
- Why AI and ML – the UAS environment will evolve rapidly, if enabled by licensed spectrum rights that enable UAS Traffic Management (UTM) that was developed by NASA. Unlike the human (air traffic controller) agents that manage and control the National Airspace System for manned aircraft, AI/ML brings the opportunity for a diverse set of AI air traffic controllers for safely managing enroute, departure, arrival, and ground control UAV operations, traffic deconfliction between manned and unmanned aircraft systems, etc. Unlike terrestrial ground vehicles (trucks and cars), open and standardized operational protocols and procedures for air traffic management have already been established for a diverse set of aircraft categories and classes, over many decades of development and continued improvement in unparalleled flight safety.

- The TAC recommends that the FCC consider serving the evolving UAS ecosystem by licensing and enabling 5G O-RAN, V-RAN, and AI/ML.

Description:

Establish a broad, comprehensive spectrum management program designed to provide the FCC with, among other things, expert analysis on the role that AI/ML can play in assisting the FCC in achieving its strategic goals. The vision incorporated in the pilot project is broadly similar to the thoughtful vision the FCC displayed when it invested substantial human and financial capital in exploring novel and largely then unproven approaches for encouraging the more efficient sharing of spectrum, culminating in CBRS. TAC recommends, as the spectrum management program's first activity, the development of a pilot project focused on the UAS use case, which would benefit from 5G O-RAN technology advancements and the application of AI/ML.

Time Frame:

The time frame for the comprehensive spectrum management program needs to be consistent with the evolving relationship over time between AI/ML techniques and communication services. Further, given the time to establish interagency policy agreements and industry engagement, the pilot program should have an initial operation period of three years.

Specific Action:

Take the steps necessary to launch the spectrum management program. This includes, among other things, identifying and securing the talent and resources needed to achieve the program's goals. Chief among the resources are talented engineers, supported by economists as well as legal and policy experts. The TAC believes that the substantial potential payoff from instituting the program warrants a significant investment in human and other capital by the FCC.

Mapping to Relevant Bureau(s):

Given the breadth of areas in which AI/ML techniques are being applied, the spectrum management program would complement the activities of the Office of Engineering and Technology (OET), the Wireless Telecommunications Bureau (WTB), the Enforcement Bureau (EB), the Public Safety and Homeland Security Bureau (PSHSB), and the International Bureau (IB).

Mapping to Strategic Objectives:

The spectrum management program has the potential to assist the FCC in reaching all of its strategic goals as identified in its strategic plan, and specifically, its interest in advancing US leadership in 5G.

Impact on the U.S.:

Due to advances in technology and substantial changes in the manner in which consumers and firms interact, participants in both the private sectors of the economy are increasing the speed with which they are adopting digital technologies. An important factor that is driving that rate of digitalization is the substantial investments that both sectors are making related to the application of AI-based techniques and methods. It is widely understood and appreciated that global leadership in AI is a necessary condition for continued growth in the U.S. economy and an increase in the standard of living in the U.S. By enhancing the ability of the FCC to make informed decisions on spectrum matters in which AI/ML plays an important role, the spectrum management program will contribute to the economic growth and prosperity of the U.S.

[1] With active RF components, interference may be modeled in the extreme and overestimated in operational environments; filtering and natural attenuation often reduces the interference. Out-of-band-emissions (OOBE) just outside of an authorized radio channel, is often the result of the non-linearity of the transmitter amplifier and practical limits of RF filtering. While transmitters are tested and radio equipment is authorized to Commission established OOBE power limits, actual emissions in operational field conditions are likely to be substantially less than the regulatory limit, because of the non-linearity of the OOBE phenomenon and the fact that most transmitters will rarely operate at maximum rated power. A little bit of lower transmit power will likely result in a ‘big’ reduction in OOBE. Hence, actual OOBE in field environments is expected to be substantially lower than worst case regulatory limits. However, lengthy disputes in Commission rulemakings often focus on theoretical worst-case emission values (or, theoretical worst-case propagation environments) that may never occur in the real world, for transmitters that are operating as designed. Interference is overestimated and spectrum may go unused for a long time.

With passive RF components, interference may not be modeled and may be underestimated in operational environments; filtering may not reduce the interference. Non-linear mixing of RF signals can lead to interference (e.g., passive intermodulation (PIM)) that is not “linearly” near the frequency/frequencies of the source of interference. A common type – 3rd order intermodulation – can result from two unrelated relatively high-power signal sources (f_1 and f_2) and cause interference at f_3 , by the relationship $2*f_1 - f_2$, or $2*f_2 - f_1$. PIM can be produced in antenna/cables of receivers, of transmitters, or many outdoor and indoor environmental structures. With the advancement and more widespread deployment of 5G infrastructure, the diversity and complexity of potential PIM sources will likely increase. The use of AI / ML may be helpful to the Commission and industry to more proactively identify, measure, and share data that can be analyzed and acted upon more quickly than traditional methods.

14 Appendix H: List of Presentations to the FCC TAC AIWG

Speaker	Affiliation	Presentation Title
Ulrika Jägare	Ericsson	“How AI is Shaping Telecom Operations”
Mazin E. Gilbert Jack Murray	AT&T Research	“AI for Autonomous Networks” (and Return Visit)
Mukarram Bin Tariq Nandita Dukkupati	Google	“Optimization of Computing and Communication Resources Using AI”
Rakesh Misra	VMware (Uhana)	“Subscriber-Centric ML/AI in Mobile Radio Access Networks” (and Return Visit)
Jason Martin	Intel (and Georgia Tech)	“Machine Learning Security & Privacy”
Berge Ayvazian	Wireless 20 20	“Breakthroughs from Synergy Between AI and 5G”
Tan F. Wong John M. Shea	University of Florida	“Dynamic Spectrum Sharing: Lessons from the DARPA Spectrum Collaboration Challenge” (and Return Visit)
Peter Volgyesi, Miklos Maroti, Peter Horvath, Sandor Szilvasi	Vanderbilt University	“Spectrum Collaboration - Building Prize-Winning Radios” (and Return Visit)
Harry Surden	U of Colorado Law School	“Artificial Intelligence, Government and Ethics”
Martin Zoltick Jennifer Maisel	Rothwell Figg	“Legal and Regulatory Considerations: Application of Artificial Intelligence to Telecommunications and the FCC” (and Return Visit)

Speaker	Affiliation	Presentation Title
Ramana Jampala	Avlino	“Predictive Modeling & Machine learning-based Optimization of Network Operations”
Jeff Alstott	IARPA	“Security of AI Systems”
Alexander Sprintson	NSF	“Impacts of AI in the Wireless Networking Domain”
Elham Tabassi	NIST - ITL	“Artificial Intelligence: A NIST strategic priority”
Rafail Ostrovsky	UCLA Stealth Software	“Stewardship of Private Data with Cryptography”
Stuart Russell	UC Berkeley	“Artificial Intelligence: History and Future”
Stephen Dennis Sridhar Kowdley	DHS S&T	“Artificial Intelligence and Machine Learning”
Ajay Vikram Singh	Nokia	“Artificial Intelligence as a Service – AI as a Service (AIaaS)”
Petros Mouchtaris	Perspecta Labs	“AI/ML Research with Applications in Telecommunications”
Thyagarajan Nandagopal Alexander Sprinston	NSF	“The NSF PAWR Initiative” (with Return Visit)
Michael G. Cotton Bradley Eales Douglas Boulware	NTIA - ITS	“EM Propagation Data, AI, and 801.22.3”

Speaker	Affiliation	Presentation Title
Kumar Navulur	DigitalGlobe - Maxar	“GIS Systems for Telecoms and AI”
Stuart Russell	UC Berkeley	“Return Visit – AI and Control”
Ken Leonard	DoT – NHTSA and RITA	“Artificial Intelligence for Transportation”
Tommaso Melodia Abhimanyu Gosain	Northeastern University	“The Use of AI in Telecommunications”
Preston Marshall	Google Wireless	“Propagation Modeling - Now Enabled by Machine Learning, Geo-Data, and Crowdsourcing”